The London School of Economics and Political Science

Essays on Peer Effects in Social Groups and Information Misperception

Manuel Staab

A thesis submitted to the Department of Economics of the London School of Economics for the degree of Doctor of Philosophy

London, July 2018

Declaration

I certify that the thesis I have presented for examination for the PhD degree of the London School of Economics and Political Science is solely my own work other than where I have clearly indicated that it is the work of others (in which case the extent of any work carried out jointly by me and any other person is clearly identified in it).

I declare that my thesis consists of approximately 51700 words.

Statement of Conjoint Work

Chapter 3, 'The Benefits of Being Misinformed', was jointly co-authored with Marcus Roel. This statement is to confirm that I contributed a minimum of 50% of this work. General decisions about the directions of research were made jointly between the authors with equal contributions to the main results.

Acknowledgements

Only the support of a great many people made this intellectual and personal journey possible.

I want to thank my supervisors whose continuous support was essential to this project. Francesco Nava has supplied an immense amount of advice, support, and constructive criticism that shaped this thesis and my entire time at LSE. And I benefited greatly from challenging, insightful, and uniquely candid conversations with Balázs Szentes. I am also grateful for the many helpful comments I received from the members of the STICERD Theory Group; especially Andrew Ellis, Erik Eyster, Gilat Levy, Matthew Levy, and Ronny Razin, who were always available for discussions. Furthermore, Christian Ghiglino offered many valuable suggestions.

My friends and colleagues made this a much more inspiring, entertaining, and certainly less solitary journey. Michel Azulai, Hao Dong, Dana Kassem, William Matcham, Davide Porcellacchia, Francesco Sannino, Arthur Seibold, Roberto Sormani, Heidi Thysen, Guo Xu, and numerous others made contributions in many ways. And I want to thank Marcus Roel who - on top of that - was also a great co-author. I furthermore want to thank Judith Shapiro for inspiring a passion for economics during my undergraduate days through her unwavering belief in the importance of economic research.

My family enabled me to pursue this journey and so I am very much indebted to them and especially to my parents and grandparents, for making this possible.

To everybody who contributed in any way: thank you!

Abstract

This thesis consists of three chapters. Chapter 1 develops a model that studies the interaction of two forces in the formation of social groups: the preference for high quality peers and the desire for status among one's peers. I present a characterization of fundamental properties of equilibrium group structures in a perfect information, simultaneous move game when group membership is priced uniformly and cannot directly depend on type. While equilibrium groups generally exhibit some form of assortative matching between individual type and peer quality, the presence of status concern reduces the potential degree of sorting and acts as a force for greater homogeneity across groups. Chapter 2 expands on this framework and analyses the effect of status concern for the provision of groups under different market structures. I particularly focus on the implications for segregation and social exclusion. I find that status concern reduces the potential for and benefit from segregation - both for a social planner and a monopolist - but the interaction of preference for rank and status can make the exclusion of some agents a second-best outcome. Social exclusion might occur even if the market is served by competitive firms.

Chapter 3 studies how two fundamental mistakes in information processing affect the welfare ranking of information experiments. In the spirit of Blackwell (1951), the binary ranking of informative action profiles is analysed under different classes of perception distortions. By themselves, an agent's tendency to misinterpret signals and the degree to which the prior deviates from the truth reduce expected utility in a model where payoff relevant actions also generate informative signals. However, experiments can be affected to different degrees. Necessary and sufficient conditions for when any binary ranking of action profiles can be reversed are provided. As a consequence, different types of mistakes can interact in non-obvious ways such that an agent might be better off suffering from both rather than just one. The chapter concludes with a characterization when such positive interaction is possible and illustrates the implications in an investment setting with costly information acquisition.

Contents

Chapter 1

The Structure of Social Groups under Status Concern

> *Because sometimes the only way you can feel good about yourself is by making someone else look bad! And I'm tired of making other people feel good about themselves.* - The Simpsons, '*Dead Putting Society*' (7F08).

1.1 Introduction

When people interact in a social environment, whether it is at work or school, in clubs or in their neighborhood, social spillovers tend to play an important role. At work, cooperation with colleagues might be essential, at school and university, studying with peers can promote understanding and enhance the learning experience. For any team sport, other players are a pre-requisite. In many of these situations, we would like to be surrounded by 'strong' peers as their ability influences the benefit we gain from the interaction. At the same time, we might want to be someone with a relatively high standing in the group. This presents a clear tension: the stronger the peers, the lower one's own standing.

Consider moving house and choosing a new neighborhood: when faced with the choice between South Kensington, one of the most affluent boroughs in London, and Camden, a borough with a more heterogeneous crowd, the decision is, among other factors, most likely influenced by the quality of public services, the valuation for these, and the price of living in the two boroughs. But in addition, one might also be worried about one's own relative status among the potential neighbors. A lower crime rate does not necessarily compensate for the discomfort caused by being one of the lowest earners.

This chapter develops a model to explore the importance of this (potential) tension in the formation of social groups very much in the spirit of Frank (1985). It addresses the question what groups can be formed and how these differ from a setting where the peer effects don't have a status component. I explore observational differences between group structures with and without status concern.

In the model, a large number of agents observe a set of prices for group membership and simultaneously decide which group to join. Agents are heterogeneous in their type: a one-dimensional variable; for example, income. The agents' payoff is determined by the composition of the group, the membership price, and their own type. In particular, two statistics of the distribution of types within a group are payoff relevant: the *quality* of the group - a function of the types of agents' choosing the same group - and the *status* of an individual - the rank in the distribution of types.[1] It is assumed that there is a positive interaction between type and the characteristics of a group. Agents with higher type value quality and rank more; just like high earners might care more about the quality of schools as well as their own social status.

It is shown that status concern reduces the possibility of segregation. Status concern is a force for homogeneity across groups as it limits the degree to which groups can differ in their composition of types. For example, if two groups are priced equally, then they have to be identical in their probability distribution over types, not just their quality. Additionally, there might be no prices that make a given group structure incentive compatible even though such prices exist if agents only care about quality. Sorting cannot be arbitrarily fine. If groups take the form of non-overlapping intervals, then the number of such intervals in equilibrium is necessarily finite. If status concern is relatively more important, less segregation can be achieved. In the extreme case where agents have preferences only over their rank, no segregation is possible and all agents joining a group pay the same price. In contrast, in the absence of status concern, arbitrarily fine sorting is possible, at least if the quality function is suitably monotone in types.

These findings can inform the literature on social groups: if status concern is relevant in a given setting, empirical investigations might lead to different conclusions and ultimately different policies. If policies are misspecified, there can be significant misallocations. More specifically, when agents care about their relative rank, we should

[1] While social status can have multiple dimension, Heffetz and Frank (2011) argue that it is inherently positional can be seen as a form of 'rank'. The simplification in this model is that agents agree on the same ranking. There is evidence that this is often the case. See Weiss and Fershtman (1998) for a survey of the relevant economic and sociological literature.

expect groups to be less segregated. If two groups are similar in quality, they should also be similar in their distribution over types. In the empirical literature on Tiebout sorting[2] - the sorting of agents into different communities based on their preferences for public goods - it is often noted that communities are much more similar across and more diverse within than should be expected.[3] This squares with the finding on segregation here. Status concern can, for a similar reason, have important implications when identifying peer effects. If we try to measure the magnitude of complementarities by the degree of segregation across groups, we need to consider how important status considerations are. An absence of positive sorting can indicate strong rank preferences rather than the absence or irrelevance of complementarities. This can, of course, lead to very different policy implications.

1.2 Related Literature

There are two themes in the literature this paper closely relates to: positional concerns and the provision of (semi)-public goods.

The notion that agents have preferences over their relative status has received considerable attention in various contexts. Veblen (1899)[4] expressed the idea of conspicuous consumption early on and Duesenberry (1949) stressed the importance of relative income in consumption and savings decisions. Generally speaking, the conspicuous consumption literature analyses the effects of preferences over consumption differences within a reference group on equilibrium outcomes. As a key contribution, the model here looks at the effects from social interaction as well as social comparison. For example, Hopkins and Kornienko (2004) study a setting where agents have preferences over their rank in a distribution of a conspicuous consumption good. Becker, Murphy and Werning (2005) characterize equilibrium income and consumption distributions when market participants can trade status. And Ghiglino and Goyal (2010) analyse how the social structure in a pure exchange economy affect equilibrium prices and allocations; finding that relatively less well-off agents can lose from social integration. Frank (1985) addresses the connection between status concern and sorting. And Ray and Robson (2012)[5] focus particularly on status as the rank in a distribution of a one-dimensional characteristic. Rank enters utility in the same way as in this model. Maccheroni, Marinacci and Rustichini (2012) provides a decision-theoretic founda-

[2]See Tiebout (1956)

[3]See, for example, Persky (1990) and Epple, Romer and Sieg (2001). Stephen Calabrese, Dennis Epple, Thomas Romer and Holger Sieg (2006) provides empirical evidence for the role of peer effects in this.

[4]See Veblen (2007) [1899]. *The Theory of the Leisure Class.* Oxford world's classics. Oxford: Oxford University Press.

[5]See also: Robson (1992).

tion for such preferences. With a stronger focus on ordinal comparisons, the literature on contests and tournaments has examined status as a way to incentivize performance. Moldovanu, Sela and Shi (2007), for example, looks at the optimal partition of agents into status categories. As a key difference to this literature, in the model here status arises automatically within a group and cannot be directly controlled by a third party.[6] Nevertheless, the models do share the zero-sum nature of status allocations.

The importance of positional concerns is also validated in the empirical literature.[7] Particular attention has been paid to the link between relative income and well-being. Alesina, di Tella and MacCulloch (2004) provide evidence that there is a significant relation between the relative income position and self-reported happiness in both the US and Europe. Card et al. (2012) exploit informational differences to investigate the role of income inequality on work satisfaction. They find that having a lower income rank than close peers has a significant negative effect on satisfaction. A similar conclusion is reached by Brown et al. (2008) where the importance of income rank is highlighted specifically. Positional concerns can also be a driver for migration decisions as shown in Stark and Taylor (1991). The role of relative income within a neighbourhood has been investigated extensively in Luttmer (2005) using data from the American Household survey. He not only finds that relative changes have an effect of similar magnitude on life satisfaction as absolute ones, but also that the effect is stronger for people that socialize more with their neighbors. Furthermore, Ashraf, Bandiera and S. Lee (2014) presents evidence from an educational setting that people are aware of their relative standing and that the salience and payoff-relevance of rank influences choices. Additionally, there is experimental evidence from Jemmott and Gonzalez (1989) that status affects performance in group settings; students performed better in groups where they have high status. Perhaps most closely related, Bottan and Perez-Truglia (2017) investigate locational preferences among medical students in the context of the National Resident Matching Program with particular focus on relative income rank. They find that people care about the cost of living and relative income rank to a similar degree. They also demonstrate that there is significant heterogeneity in the magnitude of positional concerns and that this is driven by differences in valuation of some of the spillovers generated by the different locations - in their case dating prospects.

[6] In a similar way, Rayo (2013) looks at the optimal partition of agents into status categories. See Chapter 2 for a more detailed comparison.

[7] See Frank (2005) for a brief survey of the economic literature and Weiss and Fershtman (1998) for a survey of both the economic and sociological literature.

There is a large body of literature on social spillovers and particularly the production and sharing of 'social goods' that blur the line between purely public and private goods - not unlike the quality of a group in the model proposed here. Buchanan (1965) establishes a foundation by augmenting a standard consumption model with clubs that allow the sharing of a single consumption good. Those preferences can be either directly over the sharing good or simply the characteristics of other agents. In Levy and Razin (2015), for instance, agents care directly about the average income of the people in their group. Board (2009) explores in a setting closely related to the one here the optimal group partitions offered by a monopoly supplier when these partitions also determine the social good available in each group.[8] The literature on clubs has paid particular attention to existence and stability of equilibria; primarily in the context of cooperative game theory. It has been shown that when firms providing these clubs can freely enter the market, there is a tension between heterogeneous preferences leading to differentiation and increasing returns causing greater centralization.[9] In a general equilibrium setting, Scotchmer (2005) studies the pricing of clubs with heterogenous agents. If group memberships can discriminate between relevant characteristics and thus effectively limit free movement of consumers, consumption externalities can be internalized.

The literature on networks delivers many additional insights by focusing on the specific structure within a group. Ultimately, this can be traced back to models of efficient matching, for instance Becker (1973).[10] Several papers have studied settings with social spillovers that vary across different types. In a non-specific networking environment, which is closest to the model here, agents choose how much to socialize across their network, unable to discriminate between individuals. Durieu, Haller and Solal (2011) analyse properties of Nash equilibria in such a model where agents select between discrete networking intensities that apply to all their links. Also taking the network structure as given, Bramoullé and Kranton (2007) study the provision of public goods in a network when investments can be directed to specific links. Bloch and Dutta (2009) endogenize both the link itself and the link quality. Cabrales, Calvó-Armengol and Zenou (2011) study a setting where undirected investments in links generate positive spillovers on private investment. Stable equilibria are characterized by too high or too low investment relative to the Pareto efficient outcome.

[8] See Chapter 2 for a more detailed comparison.

[9] See Aumann and Dreze (1974) for the theoretical underpinnings and Demange and Henriet (1991) for a stability result when preferences can be ordered along a single dimension.

[10] The supermodularity in traits of a match has a close analogy to the complementarity between group and individual characteristics in this model.

Baumann (2015) looks at a similar setting where benefits from (directed) investment across links are symmetric but not separable and finds asymmetric equilibria in which some agents form many low quality links while others establish few high quality connections.

The empirical literature on these social spillovers is too rich to attempt even a cursory overview here. Instead, I focus on one particular issue raised in Tiebout (1956): the endogenous sorting of agents in communities when preferences are heterogenous. Tiebout has spawned a large literature that studies the provision of public goods by competing jurisdictions that can differentiate through the public goods they offer and the taxes they charge. Theoretically, this should lead agents to cluster efficiently.[11][12] The empirical evidence, however, has been mixed. The Tiebout model in it's simple form predicts relatively homogeneous communities within i.e. fine sorting.[13] However, communities appear to be more heterogeneous within and accordingly more similar across than predicted.[14] This has been discussed in Persky (1990) and more extensively in Epple, Romer and Sieg (2001) and Calabrese et al. (2006); the latter specifically shows that this disparity can be largely resolved when allowing for preferences over the composition of the communities.[15] The model here captures some of these aspects: depending on the group quality function, there can be an incentive to separate finely but the preferences over rank can have an offsetting influence reducing the benefit from sorting and limiting the degree of segregation.

1.3 The Model

A continuum of agents drawn from the Lebesgue unit interval I choose from a countable set of groups G as part of a complete information, simultaneous move game. An agent is allowed to join at most one group. The action set is $A = G \cup \{\emptyset\}$ where $\{\emptyset\}$ represents the choice of not joining any group. Group membership is excludable through prices. For each $g \in G$, there is a charge $p_g \in \mathbb{R}^+$ for joining that group. The vector p

[11] For example, Conley and Wooders (2001) explore a settings where agents differ in tastes and genetic types. Genetic types affect the cost of acquiring a 'crowding' type that causes an externality. They characterize when taste-heterogeneous jurisdictions are optimal.

[12] Epple and Romano (2011) provides an overview of the Tiebout literature in the context of schooling choices.

[13] A prediction that has been questioned, for example in Pack and Pack (1977).

[14] See Ghiglino and Nocco (2017) for a theoretical analysis of the the interaction between conspicuous consumption and urban sorting.

[15] While Calabrese et al. (2006) delivers evidence for the importance for peer effects, the channels through which they work remain unidentified.

contains all such membership prices.[16] The price of not joining any group is normalized to 0. While these prices could be thought of more generally as costs associated with participating in a group - membership in a country club not only requires the payment of fees but also the right attire and the ability to travel there - they will, for clarity, be interpreted more literally. For example, these could be the tuition fees at a university or the membership fees of a social club.

Each agent has a one dimensional type or characteristic $w \in W$ where W is the set of characteristics in the population. W is simply taken to be the closed interval $[\underline{w}, \overline{w}]$ in $\mathbb{R}$. $\mathscr{L}$ is a strictly positive probability measure on W and $(W, \mathscr{B}_w, \mathscr{L})$ is the corresponding probability space. It admits a continuous density f with the corresponding CDF denoted F. The agent space is the atomless probability space $(\Omega, \mathscr{B}, \mathscr{P})$ where $\Omega = W \times I$ and the Borel σ-algebra $\mathscr{B} = \mathscr{B}_w \times \mathscr{B}_I$.[17] $\mathscr{P}$ is the probability measure $\mathscr{L} \times \lambda$ where λ is the Lebesgue measure on the unit interval.[18]

By joining a group, agents gain access to the peer effects generated by the agents within the same group. In particular, I distinguish between two different types of spillovers: there is a benefit ϕ that everybody values but potentially to a different degree. ϕ could be interpreted as a preference over the composition of members or simply a public good that is ‘produced’ within the group with the output level depending on the members’ characteristics.[19] I call this the *quality of the group.* Formally, ϕ is a function that maps from $\mathscr{B}$ to $\mathbb{R}^+$ and is continuous almost everywhere.[20] This can, for instance, be a statistic of the distribution of agents choosing the same group like the average type, the median type, or the lowest or highest type.[21]

Furthermore, agents have preferences over the rank in the distribution of types within their group; another statistic of the distribution of agents. I call this the *status of an agent.* Given the agents that choose g, the rank $r_g(w) = F_g(w)$ of an agent with type w is the CDF of types in g evaluated at w. In other words, rank is a function from W and probability distributions over W to $[0, 1]$.

[16]As a matter of convention, I take p to be in ascending order such that $p_i \geq p_{i-1}$

[17]$\mathscr{B}_I$ is the Borel σ-algebra of the unit interval and B_W accordingly the Borel σ-algebra of W.

[18]This follow the modelling approach for infinitely many agents suggested in He, Sun and Sun (2017).

[19]For example, Kacperczyk (2013) demonstrates the importance of university peers in the decision to become an entrepreneur - mainly through the transmission of information and a reduction in uncertainty.

[20]This is to restrict ϕ to more intuitive cases although all of the analysis would go through without this assumption given appropriate adjustment of Definition 5.

[21]Strictly speaking, it is a statistic of the distribution over $W \times A$ since the quality of a group g can depend on the measure of agents joining the group, which is not captured by the probability distribution F_g, generated by the agents choosing group g.

I characterize the agents that make the same choice in A as the *social group* - a feasible[22] measure on W.

Definition 1 (Social groups)**.** *A social group $\mathscr{L}_g$ is a Lebesgue measure on W with the property that the combination of all such social groups $\sum_{g\in A}\mathscr{L}_g(B) \leq \mathscr{L}$ for any $B \in \mathscr{B}_W$. The vector containing all social groups is denoted $\mathscr{L}_A$*

An agent of type w is in g ($w \in g$) if some agents with trait w are part of the social group $\mathscr{L}_g$ - i.e. w is in the support of that social group.

Preferences of agents are represented as follows:

$$U\left(w,g,\mathscr{L}_A\right) = u\left(w,\phi_g,r_g(w)\right) - p_g \tag{1.1}$$

Beyond additive-separability of prices, I make the following assumptions on preferences and quality:

Assumption 1 (General)**.** *$u(w,\phi r)$ is continuous and at least twice differentiable. It is strictly increasing in w, ϕ and r, and $\frac{\partial^2}{\partial w \partial \phi}u(w,\phi,r) > 0$, $\frac{\partial^2}{\partial w \partial r}u(w,\phi,r) > 0$ and $\frac{\partial^2}{\partial \phi \partial r}u(w,\phi,r) \geq 0$. For every subset $W_i \subseteq W$, ϕ is bounded for every social group $\mathscr{L}_g$ with its support contained in W_i.*

Assumption 2 (Single-crossing)**.** *If $u(\hat{w},\phi',r') \geq u(\hat{w},\phi,r)$ for some $r',r \in [0,1]$, $\phi' \geq \phi$, and $\hat{w} \in W$, then this inequality holds for all $w > \hat{w}$.*[23]

Assumption 3 (Stand-alone payoff)**.** *The stand-alone payoff for any social group $\mathscr{L}_g = 0$ as well as the payoff from the isolation choice $\emptyset$ denoted $\underline{u}$ is such that $\underline{u} \leq u(w,\underline{\phi},0)$ for all $w \in W$ where $\underline{\phi}$ is the lower bound of ϕ given W.*

The general assumptions capture the notion that agents not only value quality and rank but that this valuation is increasing in their own type. A mere scaling of W does not diminish the effect of either component. People with higher wealth living in a richer neighbourhood might exhibit stronger positional concerns. They might also have a higher valuation for school quality and other public goods exclusive to their

[22]Feasibility requires that for all measures $\mathscr{L}_g$ with $g \in A$ and all sets $B \in \mathscr{B}_W$, we have $\sum_A \mathscr{L}_g(B) \leq \mathscr{L}(B)$.

[23]In other words, u is such that the single-crossing property in the sense of Edlin and Shannon (1998) holds.

neighbourhood. [24] [25] Bayer, Ferreira and McMillan (2007), for example, provide evidence that more highly educated households value the education characteristics of their neighbours more. And Barrow (2002) finds evidence that the valuation for school quality is positively related with income and education. The fact that these preferences tend to be observable through location choice implies that they are not (entirely) obscured by status concerns. This motivates Assumption 2. It essentially states that if an agent of type w prefers higher quality over a given rank trade-off, then this must also be the case for agents with a higher type. Otherwise status concern might outweigh other preferences and impede positive sorting simply by assumption. To illustrate the assumption in terms of the university choice example, suppose a student prefers university A over university B where A offers a higher educational quality but the student's relative ability is less at A. Then, under Assumption 2, a student with higher ability would also prefer university A if she faced the same quality/rank trade-off.

Assumption 3 states that agents that choose not to participate, do not interact in this social environment. Someone who chooses not to study does not benefit from any spillovers in higher education institutions. Finally, joining any group is generally beneficial - at least at 0 cost.[26]

We can now introduce the equilibrium notion. With the next chapter in mind, I refer to the (sub-)game in which the agents observe the prices of groups and then join a group as the *agents' game*. In Chapter 2, this will be extended by an additional stage. The focus lies on pure-strategy Nash equilibria. An equilibrium is an assignment of agents to groups such that, given the social groups generated by the assignment, no agents (measure 0 set) can be made strictly better-off by an individual deviation.

Definition 2 (Equilibrium)**.** *A* ***pure Nash Equilibrium*** *in the agents' game is a $\mathscr{B}$-measurable function y from the agent space to A that either assigns each agent a group $g \in G$ or no group ($\emptyset$), with the property that for all agents $i \in \Omega$ with type $w(i)$:*

$$U\left(w(i), y(i), \mathscr{L}_A\right) \geq U\left(w(i), a, \mathscr{L}_A\right) \quad \forall a \in A$$

[24] This is similar to models where rank enters multiplicatively as in Hopkins and Kornienko (2004).

[25] This can be seen to an analogy to the 'networks as resources' view in the sociological literature. Connections with peers can be interpreted as a resource - for example due to the information peers hold. See, for example, Campbell, Marsden and Hurlbert (1986) for evidence that people with higher socioeconomic status are better connected and have thus better access to network resources. Sobel (2002) offers a critical (economic) perspective of this literature on social capital.

[26] It is, for the following analysis, without loss to simply assume $\underline{u}$ is equal to the lowest possible payoff an agent can receive in any group $\underline{u} = u(\underline{w}, \underline{\phi}, 0)$.

where $\mathcal{L}_A$ is the vector of social group generated by the assignment function y.[27]

Given an assignment, only some social groups might have a non-zero measure of agents and, mainly to avoid the negation, I call these non-empty groups the *active social groups*. An *active price* p_g is the membership price of an active social group. The smallest convex set containing the support of an active social group $\mathcal{L}_g$ is denoted $[\underline{w}_g, \overline{w}_g]$.[28] The highest type in any group g is thus $\overline{w}_g$, and the lowest $\underline{w}_g$.

In several instances the outcomes under preferences with status concern are contrasted against preferences where agents care only about quality or only about status. The utility functions below are used for these comparisons.
The utility function for agents without status concern is defined as:

$$U^q(w, g, \mathcal{L}_A) = u(w, \phi_g, r) - p_g$$

for some constant $r \in [0, 1]$. And similarly, the utility for agents with only status concern is defined as:

$$U^r(w, g, \mathcal{L}_A) = u(w, \phi, r_g(w)) - p_g$$

for some constant $\phi \in \mathbb{R}^+$.
I refer to preferences characterized by U as preferences with status concern, preferences captured by U^q as preferences without status concern, and those described by U^r as preferences over status only.

1.4 Structure of Social Groups

We start by analysing some fundamental properties of equilibrium social groups. Despite the potential multiplicity in equilibria typical for coordination games, the structure of equilibrium social groups can be characterized in terms of the relation between price, quality, and the 'extreme' types in each group. This is then contrasted to a setting with no status concern. While status concern adds some potential freedom in how distributions can overlap, it generally imposes stronger restrictions on how exactly the distributions over types can differ. This section thus explores properties that have to be fulfilled given any price vector, independent of how these prices arise. I then restrict attention to equilibria that are 'stable' in the sense that they don't rely on indifference of a continuum of agents. I exploit those results in the subsequent chapter to charac-

[27] We treat $\mathcal{L}_\emptyset$ as a 'special' social group in which agents do not interact and thus by assumption $U(w, \emptyset, \mathcal{L}_A) = \underline{u}$

[28] Clearly, for any social group $\overline{w}_g \leq \overline{w}$ and $\underline{w}_g \geq \underline{w}$

terize equilibria in a sequential game where first some provider - like a social planner or monopolist - sets these prices and then agents make their group choice.

Proposition 1. *In any equilibrium, two active social groups have equal prices $p_g = p_h$ if and only if they are identical in their probability distribution over W, meaning $F_g = F_h$.*

Proof: All omitted proofs are in the Appendix (1.6).

Let us first consider the extreme case where all prices are equal: Proposition 1 implies the probability distribution over W generated by an active social group must be identical to all other active social groups. There can be very little variation across groups. The relevant statistics - quality and status - need to be identical. Such homogeneous price vectors can arise from legal or practical restrictions. For example, in countries where tuition fees are set on a national level, this model would predict that the distribution of student ability should look very similar across universities if students care about their rank in the distribution at their university (abstracting from other factors like regional variation and locational preferences). Furthermore, this can have implications for sorting at work. If an employer lets agents freely decide on their team or choose their shift, and if employees need to be equally paid, for instance due to union rules, then we should expect the distribution of abilities within each team to reflect the overall distribution of abilities. This might give an indication why even when an employer has a strict incentive to create such representative teams, it might not be necessary to actively allocate workers across teams. [29] Without status concern, only the quality would have to be the same across active groups.

As the proposition states, this logic generalizes to any two groups that are equally priced. If they are both active groups, meaning they are chosen by agents in equilibrium, then status concern puts a strong requirement on the homogeneity across groups. Equal prices imply equal quality. But if two groups are equally priced and equal in quality, agents choose the group in which they achieve the higher rank. Therefore, the rank $r(w)$ that members with type w can attain needs to be the same across these groups. This, in turn, implies that the probability distributions over W need to be the same. The only remaining differences between groups can be in size i.e. the measure of agents in each group. Whether or not such differences can exist in equilibrium depends on the quality function ϕ. If ϕ is invariant to the size of the group and only depends on the probability distribution over types, it displays *no returns to scale* and equivalently, if ϕ changes if the measure over W is scaled, it displays *returns to scale.*

[29] See Mas and Moretti (2009)

Definition 3. *Any two social groups $\mathscr{L}_g, \mathscr{L}_h$ are identical up to size if $\mathscr{L}_g(B) = \kappa\, \mathscr{L}_h(B)$ for some $\kappa > 0$ and all $B \in \mathscr{B}_W$. They are identical if $\kappa = 1$.*

Corollary 1.1. *If in equilibrium there are two active social groups $\mathscr{L}_g, \mathscr{L}_h$ with $p_g = p_h$, then they must be such that:*
a) if ϕ has no returns to scale then $\mathscr{L}_g$ and $\mathscr{L}_h$ are identical up to size.
b) if ϕ has returns to scale then $\mathscr{L}_g$ and $\mathscr{L}_h$ are identical.

If any two active groups are priced equally, then if ϕ depends on the size of the social group, these groups must be measure-0 identical. If ϕ only depends on the probability distribution of types but not the size of the social group, then any such two social groups must be identical up-to a positive scaling parameter.

Moving to a more general price vector, we can ask how group quality, cut-off types, and prices are related. The following definition establishes a particular link between quality and extreme types. If it is fulfilled, any strict ordering in quality implies a related weak ordering in extreme types and vice versa. In its strict version, any strict ordering in quality implies a related strict ordering in type and vice versa. For instance, if social groups are strictly monotonic in the highest type, then if $\overline{w}_h > \overline{w}_g$ for two active social groups, we can also conclude that $\phi_h > \phi_g$.

Definition 4 (Monotonicity)**.** *Social groups are (weakly) monotonic in quality*
(i) at the top if for any two active groups $g, h \in G$:

$$\phi_h > \phi_g \Rightarrow \overline{w}_h \geq \overline{w}_g$$

(ii) at the bottom if for any two active groups $g, h \in G$:

$$\phi_h > \phi_g \Rightarrow \underline{w}_h \geq \underline{w}_g$$

Social groups are (weakly) monotonic in type
(i) at the top if for any two active groups $g, h \in G$:

$$\overline{w}_h > \overline{w}_g \Rightarrow \phi_h \geq \phi_g$$

(ii) at the bottom if for any two active groups $g, h \in G$:

$$\underline{w}_h > \underline{w}_g \Rightarrow \phi_h \geq \phi_g$$

I call them strictly monotonic if the inequalities are strict.

Figure 1.1: A group structure ruled out by Proposition 2

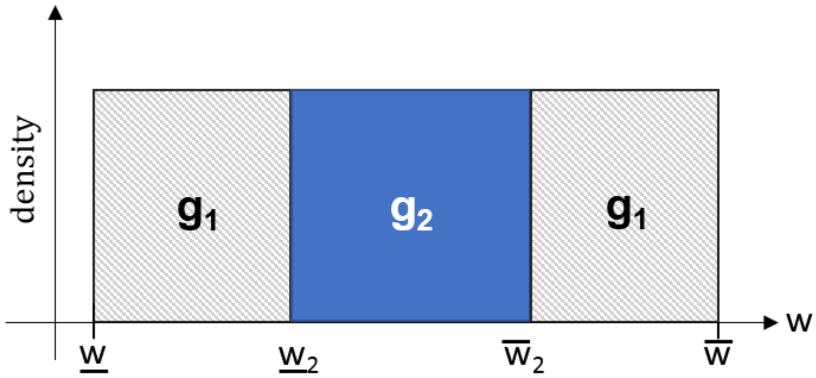

As the following Proposition establishes, we can indeed equivalently order social groups by their highest type and their quality. Furthermore, if we can strictly rank two groups by the lowest type in each group, say $\underline{w}_h > \underline{w}_g$, then the ordering in terms of quality is the same, i.e. $\phi_h > \phi_g$. And finally, if we can strictly rank two groups by their quality, say $\phi_h > \phi_g$, then this means the lowest type in h has to be at least weakly greater than the lowest type in g.

Proposition 2. *In any equilibrium, social groups are strictly monotonic in type at the top and bottom, strictly monotonic in quality at the top and weakly monotonic in quality at the bottom.*

In equilibrium, social groups are such that we can observe some degree of assortative matching between quality and type - at least at the top and bottom end of groups. The potential trade-off between rank and quality allows for a greater variety strictly inside of the support of groups. For instance, just because an agent of type w' chooses a group with quality ϕ' does not imply that all types $w > w'$ choose a group of quality $\phi \geq \phi'$. But if an agent of type w' is member in a group with quality ϕ' where she obtains rank 1, then all agents with higher type must be in a group with strictly higher quality. Intuitively, if the most able student at University A is better than the most able at University B, then University A must also have the higher educational quality.

Figure 1.1 shows a structure ruled out by Proposition 2. There are two social groups and strict monotonicity in type at the bottom implies that $\phi_2 > \phi_1$. But then strict monotonicity in quality at the top requires $\overline{w}_2 > \overline{w}_1$ which is not the case. Figure 1.2, on the other hand, shows a possible equilibrium structure with $\phi_2 > \phi_1$.

The monotonicity in quality has implications for the prices of groups in equilibrium. The larger social spillovers generated by a higher quality group require a larger mone-

Figure 1.2: A possible equilibrium structure that is not stable

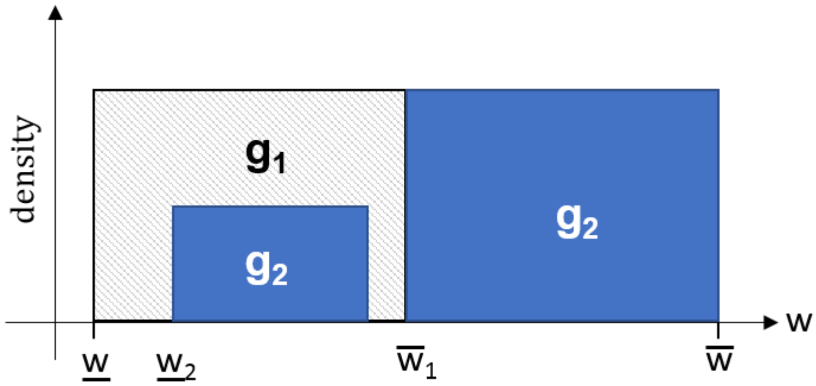

tary transfer to avoid low ranked agents in lower quality groups to join. A group with higher quality needs to have a higher price. To stay within the previous example, tuition fees at University A need to be higher. In Figure 1.2, the price of g_2 needs to be strictly greater than that of g_1.

Corollary 2.1. *In equilibrium, for any two active social groups $\mathscr{L}_h, \mathscr{L}_g$:*
$\phi_h > \phi_g \Leftrightarrow p_h > p_g$

Similar to the result that equal prices must imply equal characteristics of groups, we find that if two groups are equal in their highest type then they must be identical in their payoff-relevant characteristics:

Corollary 2.2. *For any two active social groups $\mathscr{L}_g, \mathscr{L}_h$ in equilibrium,*
(i) if ϕ has no returns to scale, $\overline{w}_h = \overline{w}_g$ if any only if the social groups are identical up to size.
(ii) if ϕ has returns to scale, $\overline{w}_h = \overline{w}_g$ if any only if the social groups are identical.

1.4.1 Stable Equilibria

Attention is now restricted to equilibria that exhibit a certain stability. If groups differ in quality, then any intersection of their supports on a set of positive measure requires the agents in this set to be exactly indifferent between higher rank versus higher quality. But indifference for a continuum of agents entails a degree of instability. Suppose, for example, a number of students over a range of abilities is indifferent between two schools. Now if the distribution of one of these schools is slightly perturbed, either because the distribution changes or because the perception of it is altered, then the status of almost all students in that school is affected. Indifference breaks necessarily for almost everybody with an ability in the overlap; even if this perturbation is arbitrarily small. When the distribution is altered by, for instance, adding a measure of agents around some w, then the rank of almost all agents above that is increased while that of those below is reduced. Independent of the effect on quality, indifference can't hold for (almost) anyone.[30]

Definition 5 (ϵ-Perturbation)**.** *A measure $\mathscr{L}_g^\epsilon$ is an ϵ-perturbation of a social group $\mathscr{L}_g$ if $\mathscr{L}_g^\epsilon \leq \mathscr{L}$ and the related quality ϕ_g^ϵ and rank r_g^ϵ differ from ϕ_g and r_g by at most ϵ.*

Clearly, for ϵ large enough, any active social group is the ϵ-perturbation of another. But for small ϵ, it allows us to describe the set of (potential) groups that are similar in status and quality. Given the continuity of ϕ, such a perturbation always exists. Stability here is the notion that if one equilibrium group was to be replaced by a very similar group (even if this group is not actually feasible), then the set of agents for which the membership in this perturbed group is suboptimal should also be small. Negligible differences in groups (or the perception of these groups) should have negligible effects on outcomes. The following refinement captures this:

Definition 6 (Stable equilibrium)**.** *An equilibrium is stable if for any ϵ-perturbation of any active social group $\mathscr{L}_g$, as $\epsilon \to 0$, the measure of agents with $w \in W$ such that*

$$U\left(w, g, \mathscr{L}'_A\right) \geq U\left(w, a, \mathscr{L}'_A\right) \quad \forall a \in A \tag{1.2}$$

where $\mathscr{L}'_A$ is the perturbed vector of social groups, approaches $\mathscr{L}_g$.

Take any equilibrium group structure. This equilibrium is stable if after the status or quality in any group is perturbed by an arbitrarily small amount, the set of agents for which the group assignment is not optimal is also arbitrarily small. As the following Lemma concludes, this rules out any overlap in between the supports of groups:

Lemma 1. *In any stable equilibrium, the intersection of the supports of all active social groups has measure 0.*

As the following result states, stability significantly narrows down the type of social groups that can form in equilibrium. At the same time, it does not pose an issue with existence. Stability simply rules out any overlaps in the support; whether complete overlaps, as in the case of social groups with equal cost that are identical up to size, or partial overlaps. The supports of all active social groups (that are not the isolation choice) form an interval partition of $[w, \overline{w}]$ where $w \geq \underline{w}$.

Proposition 3. *In a stable equilibrium, the group structure can be represented by an interval partition of $[\underline{w}_1, \overline{w}]$ with $\underline{w}_1 \geq \underline{w}$. The support of any active social group is convex.*

[30]At an endpoint, e.g. at $r = 0$, indifference might still hold but these have measure 0.

The convexity result highlights that in a stable equilibrium, there can be no 'gaps' in the support of any social group. Any such gap is driven by the rank/quality trade-off. Since the two groups can't be identical in quality, a gap in the support means that the agents with types in that gap achieve higher utility in a lower quality group.[31] But this can only arise if there are also intersections in the support. These are ruled out by stability. A group structure as in Figure 1.2 cannot be a stable equilibrium.

Stability allows for a stricter monotonicity result. Since there can be no overlaps, a higher group quality implies a higher lowest type and vice versa. Given the interval structure, we can further conclude that for any two active groups with $\phi_h > \phi_g$, we need $\underline{w}_h \geq \overline{w}_g$; the lowest type in the higher quality group needs to be weakly greater than the highest type in the lower quality group.

Corollary 3.1. *In a stable equilibrium, social groups are strictly monotonic in type and quality.*

Without the stability refinement, we were able to conclude that if two groups have equal quality, they must be equally priced and have equal support. With stability we find that there can be no two active groups of equal quality at all. It is therefore without loss to think of agents paying the same price as being members of the same group.

Corollary 3.2. *In a stable equilibrium, there can be no two active social groups equal in price or quality.*

With the combination of status concern and stability, we are able to eliminate some multiplicity. Any groups equal in quality need to be equal in their distribution over W. With stability, every type needs to be assigned just one group. Accordingly, there can be only one group of a particular quality. Following those results, not only can we represent an equilibrium group structure as an interval partition $\mathscr{I} = \{w_1, ..., \overline{w}\}$,[32] but every interval corresponds to a distinct group. In the following sections, stable equilibrium social groups will be characterized by the partition of W they induce. When convenient, I highlight the *equilibrium group structure* or equilibrium partition that a set of prices induces, rather than the social groups themselves.

Definition 7 (Equilibrium group structure)**.** $\mathscr{I} = \{w_1, ..., \overline{w}\}$ *is called an equilibrium group structure if there exists an assignment* $y : W \to A$ *such that for every interval* $[w_i, w_{i+1}]$ *with* $w_i, w_{i+1} \in \mathscr{I}$*, there is a unique* $g_i \in G$ *with* $y(w) = g_i$ *for all* $w \in (w_i, w_{i+1}]$

[31] In the opposite case, where the quality of the group covering the gap is higher, the single-crossing assumption would imply that all higher types achieve higher utility in that group. This could not be an equilibrium independent of stability.

[32] $\mathscr{I}$ is taken to be a set of points that partitions W or a subset thereof

and $y(w) = \emptyset$ for all $w \in [\underline{w}, w_1)$ and there exists a p such that this is an equilibrium in the agents' game.

1.4.2 Degree of Sorting

We now explore how status concern affects the equilibrium sorting or 'segregation'.[33] It is shown that with status concern, there is a bound on the number of active social groups in a stable equilibrium. If status concern is parametrized, it can be shown that a stronger preference for status reduces the number of active groups that can be maintained in equilibrium and, in the extreme case where only status matters, there can only be a 'representative' social group.

Without status concern, given a suitable price vector, any interval partition can be maintained in equilibrium; at least if the group quality is increasing in types.[34] The preference of agents over status in their social group puts a limit to this:

Proposition 4. *There is an upper bound $k \geq 1$ on the number of active social groups in any stable equilibrium.*

The economic intuition is that for a very fine group structure, the quality difference between any two adjacent groups is very small. But because of the interval structure, there is a cut-off type that achieves rank 1 in one group and rank 0 in the other. If groups are too close in quality, incentive compatibility for such social groups cannot hold for any vector of prices. In an equilibrium, the quality difference between two adjacent social groups needs to be sufficiently large if agents have status concern. If University A is almost identical to University B but any student in A could achieve a strictly higher rank in B, then this is not an equilibrium outcome independent of the tuition fees.

We can further show that this is 'monotone' in status-concern: as status becomes more important, the maximum number of active social groups in a stable equilibrium decreases. To capture this comparative static, I write preferences as $u_\alpha(w, \phi, r) \equiv u(w, \phi, r) + \alpha v(w, r)$ for some $0 \leq \alpha < \infty$ and a continuous and differentiable function v with $\frac{\partial}{\partial r} v(w, r) > 0$, $\frac{\partial^2}{\partial r \partial w} v(w, r) > 0$.[35] Ordering preferences by this parameter α, we can conclude that as α increases, less segregation can be maintained.

[33] I refer to a population being (unambiguously) more segregated if the equilibrium group structure is finer. This is, of course, only a partial order.

[34] If the group quality ϕ is monotone in type in the sense that if almost all types in a group are higher then the quality of that group is (weakly) higher, then any (convex) interval partition can be maintained as a group structure in equilibrium. See Definition 9 for a formal definition of monotonicity.

[35] All previous assumptions on preferences are maintained.

Corollary 4.1. *The least upper-bound on the number of active social groups in any stable equilibrium is weakly decreasing in α.*

But the fact that a stronger status concern allows for less segregation does not hinge on the stability refinement. Without stability, groups can 'overlap' and we cannot describe segregation in terms of the coarseness of the partition of W. However, in a sense groups need to be more similar. It is shown in the Appendix (1.6.2) that for any increasing sequence of α's approaching infinity, in any corresponding sequence of probability distributions of active social groups, the difference between these distributions has to converge uniformly to 0. Furthermore, as α increases, these distributions have to be arbitrarily close to the population distribution - at least over their support. If status concern becomes very important, than agents can still be excluded from participating in social groups but differences in active social groups have to disappear.

Taking this yet one step further, if agents have preferences over status only, previously defined as U^r, then no segregation can be achieved at all. If there are multiple groups, then they must be identical in their payoff relevant characteristics p and r. In a stable equilibrium, there can only be one active social group whose probability distribution is equal to the population distribution over the support of the group.[36]

Corollary 4.2. *If agents have preferences over status only, then in any equilibrium all active social groups must be identical up to size and have equal prices. In a stable equilibrium, there can be at most one active social group.*

1.5 Discussion

This Chapter establishes that status concern limits the degree of segregation across social groups that can be sustained in equilibrium when the composition of groups determines the payoff. Status concern is a force for homogeneity across groups. When people care only about the 'quality' of a group, then arbitrarily fine separation is feasible and, at least for quality functions that are increasing in the type of members, any interval partition of the type space can be sustained as an equilibrium outcome. If agents care about the rank they occupy in a group, this is no longer possible. If groups take the form of interval partitions, there is a finite upper bound on the number of non-identical social groups that can be sustained in equilibrium. In the pure status case where agents care only about their rank, there can be only one type of group.

[36] In fact, if $u_w(w,0) < 0$ where $u(w,r) \equiv u(w,\phi,r)$ then the only active social group can be the representative group, i.e. all agents join the same group.

While complementarity in types and quality allows for sorting, status concern constitutes a countervailing force.

If status concern is important, empirical models might need to take this into account. For instance, Kendall (2003) tries to derive conclusion about whether or not spillovers in ability in professional basketball exhibit complementarities - stronger players benefitting more from stronger team members - by analyzing the concentration of strong players in teams. It is found that teams are more similar in quality than would be expected if complementarities were strong. If, however, players care about their rank within the team, there might be little concentration despite the presence of complementarities. Status concern might constrain firms in how they can allocate employees across teams. In Mas and Moretti (2009) it is found that due to the nature of spillovers in ability of cashiers, it is efficient to create representative teams for each shift. Nevertheless, the firm does not actively assign workers to shifts. If these workers care about their relative ability within their team, then workers have a tendency to sort efficiently without an explicit allocation. More generally, if prices are forced to be equal across groups (e.g. laws setting tuition fees, union rules equalizing pay, etc.) then we can expect these groups to be similar in their distribution of agents' types - if agents care about status.

1.6 Appendix

1.6.1 A: Omitted Proofs

Proof of Proposition 1:

First, the following Lemma is proved:

Lemma 2. *If there are n equal active prices, in any equilibrium the corresponding n active groups must be identical in ϕ and their probability distribution over W.*

Proof. Take any probability distribution over $G \times W$. Suppose there are n active groups with equal prices. Let $G^n \subseteq G$ be the set of these groups with equal prices. Let furthermore $\overline{G}^n \subseteq G^n$ such that for any $g \in \overline{G}^n$, the corresponding social group $\mathscr{L}_g$ is such that $\phi_g = \max\{\phi_j : j \in G^n\}$.

Let $[\underline{w}^n, \overline{w}^n] = W^n \subseteq W$ be the smallest interval containing the supports of all groups in G^n. We first observe that an agent with characteristic w sufficiently close to $\overline{w}^n$ must be member of a group with $g \in \overline{G}^n$. Suppose not and she's member of a group k with $\phi_k < \phi_g$. Given equal prices, this can only be optimal if $r_k(w) > r_g(w)$.
But it follows from continuity of u that there exists an $\epsilon > 0$ such that for any $w > \tilde{w} \equiv \overline{w} - \epsilon$ we have $r_k(w) - r_g(w) < \delta$ where δ is such that:

$$u(\tilde{w}, \phi_g, r_k(\tilde{w}) - \delta) = u(\tilde{w}, \phi_k, r_k(\tilde{w}))$$

For any $w > \tilde{w}$, there is a strict incentive to join a group in $\overline{G}^n$. It follows that agents sufficiently close to $\overline{w}^n$ must be in a group with the highest quality.

Similarly, agents sufficiently close to $\underline{w}^n$ must be in a group $g \in \overline{G}^n$. We can conclude this by noting that for any other group k with $k \notin \overline{G}^n$ and $r_k(w) > r_g(w)$ for all $g \in \overline{G}^n$ there exists some $\epsilon > 0$ such that for all $w < \tilde{w} = \underline{w}^n + \epsilon$ we have $r_g(w) - r_k(w) < \delta$ where δ is such that:

$$u(\tilde{w}, \phi_g, r_k(\tilde{w}) - \delta) = u(\tilde{w}, \phi_k, r_k(\tilde{w}))$$

For an agent with a sufficiently low type, if prices are equal, there is a strict incentive to join the highest quality group since the agent's rank is close to 0 in any group and $\frac{\partial}{\partial \phi} u(w, \phi, r)$ is strictly positive for any r.

Furthermore, by the same argument, any agent in a group $k \in \{G^n \setminus \overline{G}^n\}$ with type w such that $r_k(w)$ is sufficiently close to 0 has a strict incentive to join a group in $\overline{G}^n$ since

$\frac{\partial}{\partial \phi} u(w, \phi, r) > 0$. Since every group with positive mass must have such members, all such groups need to be in $\overline{G}^n$.

To conclude, we can observe that when prices and quality are equal across a set of groups, then for any $w \in W$, we need $r_g(w) = r_h(w)$ for any $g, h \in \overline{G}^n$. All active groups need to be identical in terms of quality and rank for any type. But since the rank of an agent with type w is equal to the CDF evaluated at w, the CDF over types needs to be equal across groups. □

Using this Lemma, we can prove the Proposition:

Proof. **Sufficiency:**
If two active social groups $\mathscr{L}_g, \mathscr{L}_h$ have the same probability distribution, then necessarily $\underline{w}_g = \underline{w}_h$. For indifference, we need:

$$u(\underline{w}_h, \phi_h, 0) - p_h = u(\underline{w}_g, \phi_g, 0) - p_g$$

WLOG suppose $p_h > p_g$. Then for indifference $\phi_h > \phi_g$. But since the probability distributions are equal, we also know that $\overline{w}_h = \overline{w}_g > \underline{w}_g$. Complementarity in w, r and ϕ imply that

$$u(\overline{w}_h, \phi_h, 1) - p_h > u(\overline{w}_g, \phi_g, 1) - p_g$$

which contradicts $\overline{w}_g = \overline{w}_h$.

Necessity:
If for two active social groups $p_g = p_h$, then applying the argument from Lemma 2 to $W^n = [\underline{w}_h, \overline{w}_h] \cup [\underline{w}_g, \overline{w}_g]$ and $G^n = \{g, h\}$, we can conclude that $\phi_h = \phi_g$ which implies that for all $w \in g, h$ we need $r_g(w) = r_h(w)$ which means the probability distributions over W must be identical. □

Proof of Corollary 1.1:

Proof. a) Suppose not and for some subset $B \subset W$ we have $\mathscr{L}_g(B) \neq \kappa \mathscr{L}_h(B)$. Then it follows from the Radon-Nikodym theorem that the probability distribution over types within each group is such that for some $w \in g$, $F_g(w) \neq F_h(w)$. But it follows from Proposition 1 that this cannot be the case in equilibrium since $p_h = p_g$.
b) As the probability distributions have to be the same across both groups, suppose indeed that $\mathscr{L}_g = \kappa \mathscr{L}_h$ for some $\kappa > 0$ but $\kappa \neq 1$. If ϕ has returns to scale, then $\phi_g \neq \phi_h$ if $\kappa \neq 1$. WLOG, asumme that $\phi_h > \phi_g$. Since $r_g(w) = r_h(w)$ for all $w \in W$ and $p_h = p_h$,

all agents in g have a strict incentive to join h since they obtain the same r but higher ϕ at the same price. In equilibrium, $\phi_h = \phi_g$ and therefore $\mathscr{L}_g = \mathscr{L}_h$.

□

Proof of Proposition 2:

Proof. Take any two active social groups $\mathscr{L}_g, \mathscr{L}_h$ with with $g, h \in G$. By definition, the smallest convex set containing their support are $[\underline{w}_h, \overline{w}_h]$ and $[\underline{w}_g, \overline{w}_g]$.

Case 1 - bottom:
Suppose monotonicity in quality at the bottom fails such that $\phi_h > \phi_g$ but $\underline{w}_h < \underline{w}_g$. Then any agent with type w sufficiently close to $\underline{w}_h$ in group g can instead join group h and obtain a strictly higher rank and benefit from higher quality. This switch is strictly beneficial unless $p_h - p_g$ is sufficiently large. But due to the complementarity in ϕ and w, we know that if $\underline{w}_g > \underline{w}_h$ then

$$u(\underline{w}_g, \phi_h, 0) - u(\underline{w}_g, \phi_g, 0) > u(\underline{w}_h, \phi_h, 0) - u(\underline{w}_h, \phi_g, 0)$$

Therefore, if $p_h - p_g \geq u(\underline{w}_g, \phi_h, r_h(\underline{w}_g)) - u(\underline{w}_g, \phi_g, 0)$ then

$$u(\underline{w}_h, \phi_h, 0) - p_h < u(\underline{w}_h, \phi_g, 0) - p_g$$

Membership in h for agents close to and including $\underline{w}_h$ cannot be optimal.

Case 1 - top:
Suppose strict monotonicity in quality at the top fails such that $\phi_h > \phi_g$ but $\overline{w}_g \geq \overline{w}_h$.

Using the previous result, we know that

$$u(\underline{w}_h, \phi_h, 0) - p_h \geq u(\underline{w}_h, \phi_g, r_g(\underline{w}_h)) - p_g$$

with $r_g(w_h) \geq 0$. It follows that

$$u(\underline{w}_h, \phi_h, 0) - p_h \geq u(\underline{w}_h, \phi_g, 0) - p_g$$

From complementarity in ϕ and w, it must be the case that for all $w > \underline{w}_h$:

$$u(w, \phi_h, 0) - p_h > u(w, \phi_g, 0) - p_g$$

From complementarity between w and r, and (weak) complementarity between ϕ

and r, it follows that:

$$u(w,\phi_h,1) - p_h > u(w,\phi_g,1) - p_g$$

and therefore for any $r_g(w) \leq 1$

$$u(w,\phi_h,1) - p_h > u(w,\phi_g,r_g(w)) - p_g$$

Agents with type w such that $w \geq \overline{w}_h$ must strictly prefer membership in group h. By continuity, this is also true for agents with types below but sufficiently close to $\overline{w}_h$. A contradiction to $\overline{w}_g \geq \overline{w}_h$.

Case 2 - bottom:
Suppose strict monotonicity in type at the bottom fails such that $\underline{w}_h > \underline{w}_g$ but $\phi_g \geq \phi_h$. It must be that

$$u(\underline{w}_h,\phi_h,0) - p_h \geq u(\underline{w}_h,\phi_g,r(\underline{w}_h)) - p_g \tag{1.3}$$

Since $\phi_g \geq \phi_h$, this can only hold if $p_h < p_g$. But we know that:

$$u(\underline{w}_g,\phi_g,0) - p_g \geq u(\underline{w}_g,\phi_h,0) - p_h$$

From complementarity we conclude that this must also hold for agents with $w > \underline{w}_g$ and furthermore:

$$u(\underline{w}_g,\phi_g,r) - p_g > u(\underline{w}_g,\phi_h,0) - p_h$$

for all $r \in (0,1]$. Since $r_g(\underline{w}_h) > 0$, group membership in h cannot be optimal for agents close to and including $\underline{w}_h$. A contradiction.

Case 2 - top:
Suppose strict monotonicity in type at the top fails such that $\overline{w}_h > \overline{w}_g$ but $\phi_g \geq \phi_h$. From monotonicity in quality at the bottom, we know that $\underline{w}_g \geq \underline{w}_h$. It follows

$$u(\underline{w}_g,\phi_g,0) - p_g \geq u(\underline{w}_g,\phi_h,r_h(\underline{w}_g)) - p_h$$

This implies for any $r \in (0,1]$ and $w \geq \underline{w}_g$:

$$u(w,\phi_g,r) - p_g > u(w,\phi_h,r) - p_h$$

It follows that

$$u(\overline{w}_h,\phi_g,1) - p_g > u(\overline{w}_h,\phi_h,1) - p_h$$

A contradiction. □

Proof of Corollary 2.1:

Proof. For any two active social groups with $\phi_h > \phi_g$, monotonicity in quality at the bottom implies that $\underline{w}_h \geq \underline{w}_g$. In equilibrium:

$$u(\underline{w}_g, \phi_g, 0) - p_g \geq u(\underline{w}_g, \phi_h, 0) - p_h$$

But since for any $w \in W$, ϕ, and $r \in [0,1]$ we have $\frac{\partial}{\partial \phi} u(w, \phi, r) > 0$, we need that $p_h > p_g$.

For the other direction, suppose $p_h > p_g$. Take the type $\underline{w}_h \in h$ with $r_h(\underline{w}_h) = 0$. Necessarily, the rank obtained in the other group is $r_g(\underline{w}_h) \geq 0$. Suppose now $\phi_g \geq \phi_h$. Monotonicity in quality at the bottom implies $\underline{w}_g \geq \underline{w}_h$. In equilibrium we need:

$$u(\underline{w}_h, \phi_h, 0) - p_h \geq u(\underline{w}_h, \phi_g, 0) - p_g$$

But if $\phi_g > \phi_h$, it cannot be that $p_h > p_g$ as again $\frac{\partial}{\partial \phi} u(\underline{w}_h, \phi_h, 0) \geq 0$. And if $\phi_g = \phi_h$, it follows immediately that we cannot have $p_h > p_g$ as the above inequality would be violated. □

Proof of Corollary 2.2:

Proof. **Case (i) - sufficiency:**
It follows from strict monotonicity in quality at the top that $\phi_g = \phi_h$. Given strict monotonicity in type at the bottom, we can then conclude that $\underline{w}_g = \underline{w}_h$. This implies that $p_h = p_g$. It follows from Proposition 1 that the probability distributions need to be the same.
Case (i) - necessity:
If the CDF's are identical, then since ϕ has no returns to scale, ϕ must be identical. Strict monotonicity in type at the top then implies that $\overline{w}_h = \overline{w}_g$.
Case (ii) - sufficiency:
The argument is the same as in Case (i) with the addition that equal quality requires the measures to be identical for every measurable subset of W.
Case (ii) - necessity:
If the social groups are identical, then ϕ is identical and then strict monotonicity in type at the top implies $\overline{w}_h = \overline{w}_g$. □

Proof of Lemma 1:

Proof. Suppose there are two active social groups $\mathscr{L}_g$ and $\mathscr{L}_h$ that overlap over a set $[w_1, w_2] \subset W$. Let $\mathscr{L}_g^\epsilon$ be an ϵ-perturbation such that $\mathscr{L}_g^\epsilon$ does not differ from $\mathscr{L}_g$ over

$[w_1', w_2]$ with $w_1' > w_1$. Because of the continuity of ϕ almost everywhere, for every ϵ there exists such a perturbation (noting that rank is necessarily continuous). If for every ϵ there exists such a $\mathscr{L}_g^\epsilon$ with quality ϕ_g' such that $\epsilon > |\phi_g^\epsilon - \phi_g| > 0$, then no matter the ϵ, all agents in $[w_1', w_2]$ either strictly prefer membership in social group $\mathscr{L}_g^\epsilon$ or $\mathscr{L}_h$. The equilibrium is not stable.

If this does not exist, then we can conclude from continuity that at least for some such ϵ-perturbations, $\phi_g^\epsilon - \phi_g = 0$. But then there exist ϵ-perturbations that increase (or decrease) the status of all agents in $\mathscr{L}_g^\epsilon$ with type $[w_1', w_2]$. Again, all agents in that interval either strictly prefer $\mathscr{L}_g^\epsilon$ or $\mathscr{L}_h$ for all such perturbations. The equilibrium is not stable. □

Proof of Proposition 3:

Proof. We first prove the **interval result**:
Suppose the support of social groups does not represent an interval partition of a subset of W. Then there must be at least two social groups $\mathscr{L}_g, \mathscr{L}_h$ and a compact set $S \subset W$ such that for every subset $S_i \subseteq S$, $\mathscr{L}_g(S_i) > 0$ and $\mathscr{L}_h(S_i) > 0$. This requires that (alomst) all agents in S are indifferent between the two groups.
If not and there is a type w strictly inside S that is not indifferent and, for instance,

$$u(w, \phi_h, r_h(w)) - p_h > u(w, \phi_g, r_g(w)) - p_g$$

then for a small enough $\epsilon > 0$,

$$u(w+\epsilon, \phi_h, r_h(w+\epsilon)) - p_h > u(w+\epsilon, \phi_g, r_g(w+\epsilon)) - p_g$$

This means types immediately above w must also strictly prefer h. This implies that $r_g(w) = r_g(w')$ for all $w' > w$. Single-crossing then implies that all w' strictly prefer h to g. The support cannot intersect over S.
But if all types in S are strictly indifferent, then they have measure $\mathscr{L}(S) > 0$ which is ruled out by Lemma 1.

Finally, we observe that no set of agents $[w_1, w_2]$ strictly inside of W can be priced out of every group unless all agents with type $w < w_1$ are also not participating in any group. If there is any type $w \leq w_1$ in group g, then $U(w, g, \mathscr{L}_A) > \underline{u}$. But from single-crossing we can infer that this is also true for all $w' > w_1$.

We can now derive the **convexity result**:

Suppose the support of an active social group is not convex. Then there is an interval $W_1 = [w_l, w_h] \subset W$ such that $\mathscr{L}_g([w_l, w_h]) = 0$ for some group g and there are intervals $W_l = [\underline{w}_g, w'_l]$ and $W_h = [w'_h, \overline{w}_g]$ with $\mathscr{L}_g$ positive for any measurable subset of these two intervals and $w'_l \leq w_l < w_h \leq w'_h$.

Note that by the previous argument there must be at least one group $f \in G$ with $\mathscr{L}_f(W_1) > 0$.

The membership prices p_f and p_g must be such that for some $\hat{w} < w_h$:

$$u(\hat{w}, \phi_f, r_f(\hat{w})) - p_f \geq u(\hat{w}, \phi_g, r_g(\hat{w})) - p_g$$

If $\phi_f \geq \phi_g$, then for all $w \in (\hat{w}, w_h]$, the following must be true noting that $r_g(w)$ is constant for all $w \in W_1$ while $r_f(w)$ is increasing:

$$u(w, \phi_f, r_f(w)) - p_f > u(w, \phi_g, r_g(\hat{w})) - p_g$$

It follows from the single-crossing condition that it cannot be the case that for any $w^* \geq w_h$:

$$u(w^*, \phi_f, r_f(w^*)) - p_f \leq u(w^*, \phi_g, r_g(\hat{w})) - p_g$$

Suppose instead that $\phi_f < \phi_g$:
We know that

$$u(\underline{w}_g, \phi_g, 0) - p_g \geq u(\underline{w}_g, \phi_f, r_f(\underline{w}_g)) - p_f$$

and since u is strictly increasing in r and w, and as r_f must be constant in W_l according to the interval partition statement, we can conclude that for any $w' > \underline{w}_g$ with $w' \in W_l$:

$$u(w', \phi_g, r_g(w')) - p_g > u(w', \phi_f, r_f(w')) - p_f$$

But then it follows from the single-crossing condition that it cannot be that for any $w'' > \underline{w}_g$) that

$$u(w'', \phi_g, r_g(w'')) - p_g < u(w'', \phi_f, r_f(w)) - p_f$$

The result follows. □

Proof of Corollary 3.1:

Proof. Take any two social groups $\mathscr{L}_g, \mathscr{L}_h$. Suppose $\phi_h > \phi_g$. Weak monotonicity in quality at the bottom implies $\underline{w}_h \geq \underline{w}_g$. However, it follows from the interval result that $\underline{w}_h \neq \underline{w}_g$ and thus $\underline{w}_h > \underline{w}_g$. Everything else follows directly from 2. □

Proof of Corollary 3.2:

Proof. By definition, there can be no two groups with supports that overlap on a set of positive measure. It follows from Corollary 3.1 that groups are strictly monotonic in type which implies that if the supports do not overlap then the quality must be different.
As a direct consequence of Corollary 2.1, if two groups are not equal in quality then they can't be identical in prices. No two active social groups can be equal in price or quality in a stable equilibrium. □

Proof of Proposition 4:

Proof. Denote $\overline{\phi}$ as the upper-bound to group quality for all feasible social groups and $\underline{\phi}$ as the lower-bound. As ϕ is bounded for a given set of agents, these exist.

As in a stable equilibrium active social groups have non-overlapping, convex interval support, a necessary condition for incentive compatibility is that for any two adjacent social groups $\mathscr{L}_l, \mathscr{L}_h$ with $\phi_h > \phi_l$, we have:

$$u(\underline{w}_h, \phi_h, 0) > u(\underline{w}_h, \phi_l, 1)$$

Notice that due to the single-crossing assumption, if this inequality holds for some $\underline{w}_h$, it holds for all $w > \underline{w}_h$.
For every ϕ in $[\underline{\phi}, \overline{\phi})$ and $w \in W$, we can define $\Delta_\phi(w)$ as the minimum quality difference required such that for any $\delta < \Delta_\phi(w)$, the inequality fails for w meaning

$$u(w, \phi + \delta, 0) < u(w, \phi, 1)$$

while

$$u(w, \phi + \Delta_\phi(w), 0) = u(w, \phi, 1)$$

If no such threshold exists, which would imply

$$\lim_{\delta \to \infty} u(w, \phi + \delta, 0) < u(w, \phi, 1)$$

then set $\Delta_\phi(w) = \overline{\phi} - \underline{\phi}$.
For a given w, we can define $\underline{\Delta}_\phi(w)$ as $\inf\{\Delta_\phi(w) : \phi \in [\underline{\phi}, \overline{\phi}]\}$. Since for every $\phi \in [\underline{\phi}, \overline{\phi}]$, $\Delta_\phi(w)$ is greater than 0, this lower bound is greater than 0. Note that if $\underline{\Delta}_\phi(w)$ was not bounded away from 0, it would imply that for some ϕ^*, $\frac{\partial}{\partial r} u(w, \phi^*, r) = 0$ contradicting $\frac{\partial}{\partial r} u > 0$.

We can then define

$$\underline{\Delta}_\phi \equiv \inf\{\underline{\Delta}_\phi(w) : w \in W\}$$

as the lower bound over all these. Since for every w the difference is bounded away from 0, this is again strictly positive.

We can conclude that in equilibrium, the number of active social groups is bounded above by $\overline{N} \equiv \frac{\overline{\phi}-\underline{\phi}}{\underline{\Delta}_\phi}$. For any number of active groups $N > \overline{N}$ it must be that for at least one pair of adjacent social groups with $\phi_h > \phi_l$, we have that $\phi_h - \phi_l < \underline{\Delta}_\phi$ which implies by definition that $\phi_h - \phi_l < \underline{\Delta}_\phi(\underline{w}_h)$ which means that incentive compatibility fails. □

Proof of Corollary 4.1:

Proof. Since the number of possible active social groups in an equilibrium is a subset of $\mathbb{N}$, the least-upper-bound property implies that since it has an upper bound it has a least upper bound. Take some α as given and suppose $k_\alpha \in \mathbb{N}$ is the corresponding least-upper-bound.

Let $\mathscr{I}_{k_\alpha}$ be an equilibrium partition that has exactly k_α groups. Let $\mathscr{I}'$ be a strictly finer partition. By definition, there exists no price vector p' such that $\mathscr{I}'$ is an equilibrium partition. Let p' be the price vector such that all cut-off types $w \in int(\mathscr{I}')$ - the interior of the set - are indifferent between their group and the adjacent groups. If this does not exist, then for some $\underline{w}_h \in \mathscr{I}'$ and adjacent groups $h, l \in G$:

$$u(\underline{w}_h, \phi_h, 0) + \alpha v(\underline{w}_h, 0) < u(\underline{w}_h, \phi_l, 1) + \alpha v(\underline{w}_h, 1)$$

But if this is true for some α, it is true for all $\alpha' > \alpha$.

If instead such a p' does exist, then for some $w \in g$ and some group h with $\phi_h > \phi_g$:

$$u(w, \phi_l, r_l(w)) + \alpha v(w, r_l(w)) - p'_l < u(w, \phi_h, 0) + \alpha v(w, 0) - p'_h$$

We take h as the group immediately adjacent to l meaning that $\underline{w}_h = \overline{w}_l$ but the argument goes through for any group g with $\phi_g > \phi_l$. As α increases $\alpha(v(w, r_l(w)) - v(w, 0))$ increases. But it follows from incentive compatibility of p' at the cut-off that:

$$p'_h - p'_l = \alpha(v(\overline{w}_l, 0) - v(\overline{w}_l, 1)) + u(\overline{w}_l, \phi_h, 0) - u(\overline{w}_l, \phi_l, 1)$$

As $v(\overline{w}_l, 1) - v(\overline{w}_l, 0) > v(w, r_l(w)) - v(w, 0)$, any increase in α reduces the price difference by more than it increases the loss in utility from the lower rank. If incentive compatibility fails at some w for α, it fails for all $\alpha' > \alpha$.

Finally, note that if p' is incentive-compatible at the cut-off types, we do not need to consider agents switching from a higher to a lower-quality group because of the single-crossing assumption. Consequently, k_α is weakly decreasing in α. □

Proof of Corollary 4.2:

Proof. Suppose there are two active social groups $\mathscr{L}_g, \mathscr{L}_h$.

If prices are equal, then agents have a strict incentive to join the group in which they achieve the higher rank. In equilibrium, it must be that $r_g(w) = r_h(w)$ for all $w \in W$ which implies that the CDF's are identical.

Now suppose prices are not equal and assume wlog that $p_h > p_g$. We can deduce that agents close to $\overline{w}$ must be in g since the rank difference between groups is arbitrarily close to 0 for agents sufficiently close to $\overline{w}$. Furthermore, there exists an $\epsilon > 0$ such that agents with types in $[\underline{w}_h, \underline{w}_h + \epsilon)$ must also be in g. This follows since $r_g(\underline{w}_h) \geq r_h(\underline{w}_h) = 0$. For a small enough ϵ, all agents in $[\underline{w}_h, \underline{w}_h + \epsilon)$ have a strict incentive to be in g since the rank difference $r_g(w) - r_h(w)$ is either positive or sufficiently small. But this is a contradiction since there exists a $\delta > 0$ such that agents in $[\underline{w}_h, \underline{w}_h + \delta)$ must be in h by definition of $\underline{w}_h$. It must be that $p_h = p_g$. In a stable equilibrium, this cannot be the case which follows from Corollary 3.2. The result follows. □

1.6.2 B: Convergence

The following result shows that, loosely speaking, the stronger the status concern, the more similar any active groups have to be in equilibrium. Given a parametrized specification of preferences, for a given α_n let $\mathscr{L}_g^n$ and $\mathscr{L}_h^n$ be any two equilibrium social groups for some price vector p. There always exists such an equilibrium since $\mathscr{L}_g$ and $\mathscr{L}_h$ can be identical and equally priced. We can denote the corresponding probability distributions over W as F_g^n and F_h^n. The following result shows that for any sequence of $(\alpha_n)_{n\in\mathbb{N}}$ that is strictly increasing and converging to infinity, any corresponding sequence $(G^n)_{n\in\mathbb{N}}$ where $G^n \equiv F_g^n - F_h^n$ - the difference of two equilibrium probability distributions given α_n - must converge to 0 and this convergence is uniform.

Proposition 5. *For any increasing sequence $(\alpha_n)_{n\in\mathbb{N}}$ converging to infinity, the difference of the probability distributions over W for any corresponding sequence of two active social groups in equilibrium converge uniformly to 0.*

Proof. For any two groups h, l with $\phi_h \geq \phi_l$, we can establish an upper bound on the rank difference such that there can exist prices that make these two groups incentive compatible. As the price of any group must be increasing in ϕ (see Corollary 2.1), we know that $p_h \geq p_l$ with strict inequality if the quality difference is strict.

For every $w \in W$, let

$$\overline{\Delta}_r(w) = \sup R_w$$

where

$$R_w = \left\{r_l - r_h : r_l, r_h \in [0,1] \,\&\, u(w,\phi_h,r_h) - u(w,\phi_l,r_l) + \alpha\left[v(w,r_h) - v(w,r_l)\right] > 0\right\}$$

This is the maximum rank difference that can be sustained in equilibrium for any w. If $\overline{\Delta}_r(w) = 0$, then it must be that $\phi_h = \phi_l$ or there can be no such two groups in equilibrium as incentive compatibility requires that $\overline{\Delta}_r(w) > 0$ otherwise $p_h - p_l \leq 0$ which implies that incentive compatibility fails for type w if $\phi_h > \phi_l$.

As ϕ is bounded, we can conclude that this maximum rank difference is bounded for every w by:

$$\overline{\Delta}_r^*(w) = \sup \overline{R}_w$$

where

$$\overline{R}_w = \left\{r_l - r_h : r_l, r_h \in [0,1] \,\&\, u(w,\overline{\phi},r_h) - u(w,\underline{\phi},r_l) + \alpha\left[v(w,r_h) - v(w,r_l)\right] > 0\right\}$$

We can define the upper bound over all these $\overline{\Delta}_r^*(w)$ for a given α as

$$\overline{\Delta}_r^*(\alpha) = \max\{\overline{R}_w : w \in W\}$$

If α is sufficiently large, then due to the boundedness of utility we have $\overline{\Delta}_r^*(\alpha) < 1$. This implies that $F_h(w) - F_l(w) \leq \overline{\Delta}_r^*(\alpha)$ for all $w \in W$.

Take any increasing sequence $\{\alpha_n\}_{n\in\mathbb{N}}$ where $\alpha_n \to \infty$ as $n \to \infty$.

For every $\epsilon > 0$, there exists an N such that $\overline{\Delta}_r^*(\alpha_n) < \epsilon$ for all $n > N$. This means that $F_h - F_l$ converges uniformly to 0. As $\overline{\phi}$ and $\underline{\phi}$ are upper and lower bounds on the quality of any two groups and as $u_\phi > 0$, the maximum difference in rank for any two qualities $\overline{\phi} > \phi_h \geq \phi_l \geq \underline{\phi}$ is bounded by $\overline{\Delta}_r^*(\alpha)$. The result follows. □

Corollary 5.1. *For any increasing sequence $(\alpha_n)_{n\in\mathbb{N}}$ converging to infinity, the probability distribution over $(w_1, \overline{w}] \subseteq W$ for some $w_1 \geq \underline{w}$ for any corresponding sequence of active social groups in equilibrium converges uniformly to the population distribution F.*

Proof. If in any equilibrium there is an active social group $\mathscr{L}_g$ that somewhere over its support differs from $\mathscr{L}$ on some subset $W_d \subseteq W$, then it follows from Assumptions 1 and 3 that there must be another active social group $\mathscr{L}_h$ different from $\mathscr{L}_g$ with $\mathscr{L}_h(W_d) > 0$. But it follows from Proposition 5 that if there are two active social groups, then as $\alpha \to \infty$, they must be arbitrarily similar. But then either there is only one active group in some interval, in which case $\mathscr{L}_g = \mathscr{L}$ over its support. Or there is more than one but as $\alpha_n \to \infty$, the difference in their probability distributions must go to 0 and so they both must converge to $\mathscr{L}$ which implies that their associate probability distributions converge to F. It again follows from 5 that this convergence is uniform. □

Chapter 2

The Provision of Social Groups under Status Concern

2.1 Introduction

This chapter extends the previous analysis and explores the consequences of status concern on aggregate welfare and revenue. Rather than just asking what groups can be offered, we might be interested in what groups a social planner maximizing utilitarian welfare or a private firm maximizing profit want to offer. And in particular, this chapter tries to answer how these optimal provisions are affected by the presence of status concern. The focus lies on two key aspects: segregation and social exclusion. It is explored how status concern affects the segregation of agents; i.e. how fine is the optimal group structure offered by a social planner, monopolist, or in an extension, competitive firm. And it is examined what status concern implies about social exclusion, addressing the question of how many agents might not be offered any social group.

While Chapter 1 examined properties of social groups that hold for any set of prices, this Chapter extends the setting by introducing an additional stage: a seller or 'provider' posts a menu of prices for social groups and then agents choose from this set of prices. This is modelled as a direct communication mechanism where agents report their type and the mechanism recommends a price and group. For simplicity, it is assumed that offering groups is costless. This provider could, for example, be a local authority deciding on the number, type, and tuition fees of schools in the district; or a firm developing a new housing project, choosing how inclusive the development should be. The provider might act as a benevolent social planner - the authority maximizing aggregate welfare - or as a monopolist - the authority maximizing profits. The role of competition is also discussed.

It is shown that, in addition to the reduced potential for segregation which was observed in the previous chapter, status concern also reduces the benefit from segregation. More precisely, 'splitting' a population into several, separate groups is less beneficial under status concern - both in terms of aggregate welfare and, under some restrictions, in terms of revenue. Status concern not only leaves less room to manoeuvre in terms of which groups can be offered, a provider also gains less from offering a finer group structure. In the extreme case where agents only care about status, the unique welfare maximizing group structure is the representative group in which all agents participate. In contrast, Board (2009) finds in a closely related setting without status concern that for sufficiently convex quality functions, full separation can indeed be both a welfare and profit maximizing equilibrium.

As a second key observation, the interaction between quality and status concern can, in combination with the limitation to anonymous pricing, make the exclusion of some agents from any social group a second-best outcome. This is true even if in the first-best, where agents can be directly assigned a group, full participation is optimal. If agents care only about quality or only about status, this cannot be the case. In this sense, the agents' concern for status *and* quality can lead to social exclusion. The group provider achieves social exclusion by setting all prices high enough so that some agents strictly prefer not to join any group. Social exclusion can only be maintained at the 'bottom'; the set of excluded agents forms an interval at the low-end of the type distribution. In the context of the education example, even an authority maximizing utilitarian welfare might set university tuition fees such that some students choose not to acquire higher education. If we alter the population distribution - suppose a group of new agents arrive in a society - the planner might want to raise prices in order to exclude the low-type arrivals even though their utility enters the planner's objective function. Similarly, if a new group of high-type agents arrives, a planner might set prices such as to exclude some low-type agents that were previously members of a group. Maybe surprisingly, in some cases a monopolist might charge a lower price for the lowest-quality group and thus exclude fewer agents than a social planner. The presence of status concern can imply that the welfare effects of otherwise unambiguous policy interventions become less straight-forward. If, for instance, a policy maker aims to reduce social exclusion, this might come at a cost of lower aggregate welfare.

2.2 Related Literature

The model presented here draws from two closely related theoretical papers: Board (2009) investigates the optimal monopoly pricing of social groups when the agents' types determine the quality of the group.[1] It is shown that independent of the exact nature of the quality function, the monopolist provision is too segregated and excludes too many agents. As key distinction, in Board (2009) agents value the distribution of types only in terms of group quality; the payoff is independent of other aspects of the underlying distribution of types. In Rayo (2013), the agents' type is their private information and they obtain status through signals sold by a monopolist. The monopolist thus controls the agents' status. This can lead to pooling for some subsets of agents and full-separation for others. Broadly speaking, in Board (2009) agents care about local quality - the quality of their social group - while in Rayo (2013) they care about global status - the status in the population. In the model developed here, agents have preferences over local quality and local status.

In a broader sense, the model here is also related to a large body of literature on monopoly screening and product variety. For example, Brito and Oakland (1980) examine the optimal monopoly provision of excludable public goods, finding that a monopolist underprovides goods. The setting here differs as agents not just care about their 'product category' but also about the other agents in the same category. Mussa and Rosen (1978) examines the problem of a monopolist differentiating products to discriminate between different consumers. In their model, a monopolist generally underprovides quality, overly differentiates products, and inefficiently excludes agents. While all of these aspects can be present here, it is demonstrated that status concern can lead to too little sorting and exclusion compared to the social planner solution.

2.3 Setting

The set-up here closely follows the model introduced in Chapter 1. A group provider offers a menu of prices and groups. The agents, after observing this menu, choose which group to join. This is modelled as a direct communication mechanism where agents report their type to the mechanism (the message space being restricted to W) and then get assigned a group and payment. The solution concept is a correlated equilibrium (in the sense of Myerson (1982) and Aumann (1987)). As a key assumption, the

[1]Windsteiger (2017) presents a sorting model where agents care about the average type in their group and a monopolist can offer two segregated groups, or a representative group. It is shown that the monopolist's incentives do not generally align with the socially optimal outcome.

mechanism designer has no direct control over the status of an individual. If an agent was to submit a different type, it is assumed that agents within the same group still recognize the agent's true type and he thus obtains the same status. In contrast to Rayo (2013), it is not the mechanism that awards status. An agent's type is known by the other members without any involvement of the designer. This captures settings where agents have good information about their peers or can easily signal their type.[2]

Definition 8 (Group provision)**.** *A group provision is a $\mathscr{B}_W$-measurable function*

$$m : W \to A \times \mathbb{R}^+$$

This group provision generates social groups in the sense of Definition 1 (Chapter 1). For every $g \in G$, the associated social group $\mathscr{L}_g$ is such that for every $B \subset W$, $\mathscr{L}_g(B) = \mathscr{L}(m^{-1}(g \times \mathbb{R}^+) \cap B)$. The resulting vector of social groups is again denoted $\mathscr{L}_A$.[3]

Analogous to the notation in Chapter 1, the group assignment of an agent of type w consistent with m is denoted $y(w)$ and the price as $p(w)$ so that the group provision can be decomposed into $m(w) = (y(w), p(w))$. Agents' preferences follow the same assumptions as in Chapter 1 and are characterized as follows:

$$U(w, m(w'), \mathscr{L}_A) = u(w, \phi_{y(w')}, r_{y(w')}(w)) - p(w') \tag{2.1}$$

where $w' \in W$ is the type reported by the agent. Note that even if w' is reported, the agent still obtains rank $r_{y(w')}(w)$ where $y(w')$ is the group assignment following report w'.

The planner problem can be written as a problem of setting prices and assigning agents to groups such that the group provision is incentive compatible and individually rational (i.e. satisfies the participation constraint). Incentive compatibility requires that reporting the true type is optimal for all agents:

$$U(w, m(w), \mathscr{L}_A) \geq U(w, m(w'), \mathscr{L}_A) \quad \forall w, w' \in W \tag{2.2}$$

We can then state the planner problem as:

[2]Ashraf, Bandiera and S. Lee (2014) provides evidence in an educational setting where this is indeed the case. Participants in a training scheme seem to be aware of their position in the distribution of relative abilities.

[3]While the stand-alone choice $\emptyset$ is included in the group provision, individual rationality requires the price to be 0

$$
\begin{aligned}
&\max_{m(w)} \int_W \big[U(w,m(w),\mathscr{L}_A) + p(w)\big]\,dF(w) \\
&s.t. \quad U(w,m(w),\mathscr{L}_A) \geq U(w,m(w'),\mathscr{L}_A) \quad \forall w,w' \in W \\
&\qquad\;\; U(w,m(w),\mathscr{L}_A) \geq \underline{u} \quad \forall w \in W
\end{aligned}
\tag{2.3}
$$

where F is the population distribution over W as defined by $\mathscr{L}$, and $\mathscr{L}_A$ is the vector of social groups generated by $m(w)$. Incentive compatibility then requires that $p(w)$ is constant over the support of each group: [4]

Lemma 3. *Every incentive compatible group provision $m(w) = \big(y(w), p(w)\big)$ is such that $p(w) = p(w')$ for all $w, w' \in W$ with $y(w) = y(w')$.*

Proof. All omitted proofs are in the Appendix (2.8). □

Following this result, we can write the planner problem as the optimal choice of assignment $y(w)$ and membership prices p:

$$
\begin{aligned}
&\max_{y(w),p} \int_W \big[U(w,m(w),\mathscr{L}_A) + p_{y(w)}\big]\,dF(w) \\
&s.t. \quad U(w,m(w),\mathscr{L}_A) \geq U(w,m(w'),\mathscr{L}_A) \quad \forall w' \in W \\
&\qquad\;\; U(w,m(w),\mathscr{L}_A) \geq \underline{u} \quad \forall w \in W
\end{aligned}
\tag{2.4}
$$

Since incentive compatibility requires the planner to offer a uniform price for each group, we can as in Chapter 1 denote the vector of all group prices as p. As a convention, this does not include the 'price' for the stand-alone social group $\mathscr{L}_\emptyset$ which is by default 0. As before, 'active' groups refers to social groups chosen by a positive measure of agents. Despite the change in focus, the results from the previous chapter still hold, as Proposition 6 shows:

Proposition 6. *A group provision $m(w) = \big(y(w), p(w)\big)$ is incentive compatible and individually rational if and only if $p(w)$ is constant for each active social group and $y(w)$ is a stable equilibrium in the agents' (sub-)game given p.*

We can conclude that the optimal group provision is equivalent to the planner-optimal, stable equilibrium in the agents' game. This allows us to proceed with the notation from Chapter 1 (as well as the relevant results) despite the shifted focus and modified setting.

[4] As $y(w)$ is a function, the supports of the social groups generated by the corresponding $m(w)$ intersect on a measure 0 set. We can thus again describe the social groups in terms of the partition $\mathscr{I}$ of W they generate.

2.4 Status and Welfare

This section closely examines the problem of a social planner maximizing aggregate welfare. We can think of the planner as a local government authority planning the provision of public goods that involve some form of social interaction. In the context of education, the local authority in a school district might determine which schools to offer and how to set entry barriers. Or a housing board might want to plan new housing developments to maximize aggregate welfare. The question here is whether to set prices such as to segregate residents according to some dimension, e.g. income, or create a more inclusive housing project. Any transfer paid by an agent enters the objective function and is part of aggregate welfare (as opposed to money burning).[5] Since the transfers cancel out, the planner maximizes welfare as if there were none. Prices simply serve the purpose of maintaining incentive compatibility:

Incentive Compatibility and Prices

Before developing results on segregation and exclusion, let us first examine the constraint incentive compatibility poses on the groups that a planner can provide.

For each active social group $\mathscr{L}_g$ and $w \in g$, define:

$$\underline{\Delta}_r(w,g) \equiv u(w,\phi_g,r_g(w)) - u(w,\phi_g,0)$$

and

$$\overline{\Delta}_r(w,g) \equiv u(w,\phi_g,r_g(w)) - u(w,\phi_g,1)$$

This is the difference between the actual utility obtained by a w-agent and the utility he would receive if he had rank 0 and 1 respectively. Take any two adjacent groups g and h with $\phi_h > \phi_g$. Incentive compatibility holds if and only if prices are such that for all $w \in g$:

$$u(w,\phi_h,0) - u(w,\phi_g,0) - \underline{\Delta}_r(w,g) \leq p_h - p_g \tag{2.5}$$

and for all $w \in h$:

$$u(w,\phi_h,1) - u(w,\phi_g,1) + \overline{\Delta}_r(w,h) \geq p_h - p_g \tag{2.6}$$

At the boundary between any two active social groups $\mathscr{L}_h$ and $\mathscr{L}_g$, we have a cut-off type $\hat{w} = \underline{w}_h = \overline{w}_g$ with $r_h(\hat{w}) = 0$ and $r_g(\hat{w}) = 1$. Since agents of this type need to be

[5] Recall that $U(w,m(w),\mathscr{L}_A) \equiv u(w,\phi_{y(w)},r_{y(w)}(w)) - p(w)$.

indifferent between both groups, incentive compatibility completely pins down the price difference:

$$p_h - p_g \geq u(\hat{w}, \phi_h, 0) - u(\hat{w}, \phi_g, 0) - \underline{\Delta}_r(\hat{w}, g)$$
$$p_h - p_g \leq u(\hat{w}, \phi_h, 1) - u(\hat{w}, \phi_g, 1) + \overline{\Delta}_r(\hat{w}, h)$$

which imply

$$p_h - p_g = u(\hat{w}, \phi_h, 0) - u(\hat{w}, \phi_g, 1) \tag{2.7}$$

Without status concern, $\overline{\Delta}_r(w, h) = \underline{\Delta}_r(w, g) = 0$ for all w. Given the complementarity between type and quality, incentive compatibility is thus satisfied for all agents if (2.7) holds for any two adjacent groups. Under status concern, this is not necessarily the case. The single-crossing assumption guarantees that if (2.6) is satisfied for some type w, it also holds (strictly) for all $w' > w$. However, 'upward' incentive compatibility in (2.5) might not hold. Problems with incentive compatibility can arise because $\underline{\Delta}_r(w, g)$ is increasing in w. An agent with low rank in group g loses less from switching to a higher quality group because of the lower rank in g.[6] The smaller the quality difference $\phi_h - \phi_g$, the more of a problem this is. Furthermore, if the quality difference is 'too small', then no prices might exist such that (2.7) holds.

Suppose a local school authority tries to separate students along different ability levels. If the quality of education is nevertheless similar across these schools, then such a separation cannot be achieved through tuition fees. This can arise for two reasons: students that have a low rank in the lower quality school might strictly prefer the higher quality school given the price difference pinned down by the cut-off type. Or low ranked students in the higher quality school strictly prefer to switch down and achieve a high rank. There is no set of prices that achieves the desired cut-off.

This discussion leads to two observations: First, a social planner can only provide social groups such that given the dispersion in types, the quality differences between groups are sufficiently large. And second, since only price differences matter, any incentive compatible full-participation group structure can be offered with the lowest price 0. For an equilibrium provision, we can thus focus on incentive compatibility alone as the participation constraint poses no issue.

[6] Stark and Taylor (1991) present evidence that a low relative income rank increases the probability of (temporary) migration to a higher income country relative to internal migration. This is consistent with the notion that people with a lower rank have less status to lose when switching to a new reference group. They can have a stronger incentive to join to a group with higher spillovers despite the complementarities in type and quality.

2.4.1 Status and Segregation

Without status concern, the degree of segregation - or sorting - a social planner would want to induce is determined by the 'convexity' of the group quality. If splitting a group leads to a sufficient increase in average quality, such a split is welfare improving. This can lead to arbitrarily fine sorting. For example, if the quality of a group only depends on the lowest type, then any split increases aggregate welfare. As was shown in Proposition 4, under status concern the number of groups in any equilibrium is necessarily finite. Status concern imposes a limit to the degree of segregation that can be sustained in equilibrium. As will be shown in this section, status concern also lowers a planner's incentive to segregate in the first place - at least as long as status concern does not change the valuation of quality on aggregate.

General welfare comparisons in this context are difficult to make without imposing strong assumptions on how preferences with and without status concern relate to each other. To minimize these, I restrict attention to Welfare rankings of group structures that can be unambiguously compared in terms of coarseness - i.e. segregation - as well as participation - i.e. social exclusion.

The following assumption rules out that the effect of changes in quality on aggregate depends on status concern. If we look at an entire group, then under Assumption 4, on average agents value quality the same whether or not they have preferences over status. For example, if we add up all the individual valuations within a school, then a change in school quality should have the same aggregate welfare consequences whether or not students care about their relative standing within the school. A sufficient (but not necessary) condition for this to be fulfilled is that preferences are separable in quality and status - thus closely resembling the utility-representation derived in Maccheroni, Marinacci and Rustichini (2012). This assumption for a cleaner result but an equivalent statement could be made as long as this interaction is sufficiently small.

Assumption 4 (Status-quality neutrality)**.** *Preferences are such that for every measure* $\mathscr{L}_g \leq \mathscr{L}$ *with support over* $[w_1, w_2] \subseteq W$ *and associated* ϕ *and* r*:*

$$\frac{\partial}{\partial \phi} \int_{w_1}^{w_2} u(w, \phi, r(w)) dF(w) = \frac{\partial}{\partial \phi} \int_{w_1}^{w_2} u(w, \phi, r_0) dF(w)$$

where $r_0 \in [0,1]$ *is the constant rank for* $U^q(w, m(w), \mathscr{L}_A) \equiv u(w, \phi_{y(w)}, r_0) - p(w)$.

It is now shown that given any incentive compatible group provision, offering a finer group structure has a less positive (more negative) effect on welfare if agents have preferences over rank than with preferences over quality only.

Suppose there exist group provisions $m(w)$ and $m^q(w)$ for preferences with (U) and without (U^q) status concern such that the partition $\mathscr{I}$ represents the associated equilibrium group structure. Suppose further that $\mathscr{I}'$ is a strictly finer partition than $\mathscr{I}$ and there are also incentive compatible provisions $m'(w)$ and $m^{q\prime}(w)$ that induce $\mathscr{I}'$ under the respective preferences.

Proposition 7. *Under Assumption 4, offering $m'(w)$ compared to $m(w)$ achieves higher aggregate welfare under status concern only if offering $m^{q\prime}(w)$ compared to $m^q(w)$ achieves higher aggregate welfare without status concern.*

Intuitively, if we could freely allocate ranks to agents, the complementarity between type and rank would lead us to assign high ranks to high-type agents. If any social group is split, then almost all agents in the higher group are assigned a lower rank than before and almost all agents in the lower group a higher rank. Since the total measure does not change, this leads to a drop in welfare due to the complementarity between type and rank. Since the effect of a change in ϕ is the same with and without status concern, given this re-assignment of ranks, we can conclude that the overall effect on welfare is less positive under status concern. Segregation (potentially) helps to match quality and type efficiently but there is a necessary loss from the mismatch between type and rank.

This tension between the 'allocation' of status and quality can be further illustrated by again parametrizing status concern. As status becomes relatively more important, the benefit from segregation falls. To capture this comparative static, we write preferences again as $u_\alpha(w,\phi,r) \equiv u(w,\phi,r)+\alpha v(w,r)$ for some $0 \le \alpha < \infty$. Ordering preferences by this parameter α, we can conclude that as α increases, segregation not only becomes harder to maintain but is also less beneficial.

Let $m_\alpha(w)$ and $m_\alpha{}'(w)$ refer to incentive compatible group provisions that, under status concern with parameter α, lead to the group partitions $\mathscr{I}$ and $\mathscr{I}'$, with $\mathscr{I}'$ finer than $\mathscr{I}$:

Corollary 7.1. *With preferences $u_\alpha(w,\phi,r)$, the group provision $m_\alpha{}'(w)$ achieves higher aggregate welfare than $m_\alpha(w)$ only if $m'_{\hat{\alpha}}(w)$ achieves higher aggregate welfare than $m_{\hat{\alpha}}(w)$ under $u_{\hat{\alpha}}(w,\phi,r)$ for all $\hat{\alpha} \in [0,\alpha]$.*

In the extreme case where agents only have preferences over status, segregation into several social groups cannot be maintained and is not optimal. If so, the second-best outcome is unique (up to the level of transfers) and is equal to the first-best.

Corollary 7.2. *With preferences over status only (U^r), the full participation group $\mathscr{L}_g = \mathscr{L}$ is the unique welfare-maximizing group structure. There exists an equilibrium group provision $m(w)$ that induces this group structure.*

Example: Segregation and welfare loss

While status concern makes segregation less beneficial, it can nevertheless cause a large welfare loss if the reduced possibility to segregate agents interacts sufficiently negatively with the group quality. This example demonstrates such a setting.

Consider a quality function ϕ that in sensitive to 'congestion'. If a group exceeds a certain size, the benefit to the members decreases:

$$\phi_g = \overline{w}_g + \mathbb{1}_{\mathscr{L}_g(W)\leq\epsilon}(K_1 - K_0) + K_0$$

where $K_1 > K_0 > 0$, $\epsilon \in (0,1)$, and $\mathbb{1}$ the indicator function. For a group with measure less or equal ϵ, the quality equals the highest type plus a constant K_1. If the size of the group surpasses the threshold ϵ, then the quality of the group drops by $K_1 - K_0$. For simplicity, let types be distributed uniformly over $[\underline{w}, \overline{w}]$ and the utility be separable in rank and quality:

$$u(w,\phi,r) = w\phi + \alpha w(r - \frac{1}{2})$$

As the quality of a given group is determined (beyond the measure) by the highest type, if there was no congestion ($\epsilon = 1$), the welfare maximizing partition would be the representative group. Any further segregation lowers the quality of some groups, without increasing the quality of any other group. Due to the congestion, the representative group is no longer optimal if $K_1 - K_0$ is large enough. In this case, some segregation is optimal. In particular - ignoring incentive compatibility - if the increase in quality by avoiding the negative effect of congestion is large enough, then in the optimal provision groups are segregated as much as necessary and are all reduced to measure less or equal ϵ. Without status concern, there exist prices such that this can be achieved in an incentive compatible way. As the following two cases show, a relatively small increase in status concern can, however, lead to a large welfare loss if it means that the first-best group structure can no longer be implemented:

Case 1: $\alpha = 0$, $K_1 - K_0 > \overline{w} - \underline{w}$
Agents don't exhibit status concern. The welfare-maximizing group structure is the interval partition I^ϵ with $F(w_i) - F(w_{i-1}) = \epsilon$ for all $w_i \in I^\epsilon$ with $i > 2$. All groups have exactly measure ϵ (except for the lowest quality group which has measure $\leq \epsilon$). Because of the complementarity between w and ϕ, and as ϕ_i is strictly increasing in w_i, prices exist such that I^ϵ is an equilibrium group structure.

Let $\epsilon' \equiv \frac{\epsilon}{\overline{w}-\underline{w}}$. Total surplus is bounded below by:

$$\int_{\overline{w}-\epsilon'}^{\overline{w}} w(K_1 + \overline{w})dw + \int_{\underline{w}}^{\overline{w}-\epsilon'} w(K_1 + \underline{w})dw \tag{2.8}$$

Case 2: $\alpha = \epsilon'$, $K_1 - K_0 > \overline{w} - \underline{w}$
In this case, agent's exhibit status-concern. With $\alpha = \epsilon'$, there exist no prices that implement the group structure I^ϵ in an incentive compatible way as for any $w_i \in I^\epsilon$ at the boundary between g_i and g_{i+1}:

$$u(w_i, \phi_{i+1}, 0) - u(w_i, \phi_i, 1) = w_i(w_i + \epsilon' - w_i) + \epsilon' w_i(0-1) = 0$$

which cannot be incentive compatible (see Corollary 2.1). Active groups (except for the highest-quality one), must have measure greater than ϵ. Maximum social surplus is then bounded above by

$$\int_{\overline{w}-\epsilon'}^{\overline{w}} w(K_1 + \overline{w})dw + \int_{\underline{w}}^{\overline{w}-\epsilon'} w(K_0 + \overline{w})dw + \frac{1}{8}\epsilon'(\overline{w} - \underline{w}) \tag{2.9}$$

Comparing Cases 1 and 2, the difference in welfare is at least:

$$\frac{1}{8}\epsilon' - \int_{\underline{w}}^{\overline{w}-\epsilon'} w(K_1 - K_0) + \int_{\underline{w}}^{\overline{w}-\epsilon'} w(\overline{w} - \underline{w}) \tag{2.10}$$

As $(K_1 - K_0) \to \infty$, the welfare loss from the ϵ' change in status-concern becomes arbitrarily large. Without status concern, the complementarity between quality and type allows for enough flexibility in providing groups that congestion can be avoided. With status concern, this might no longer be possible causing a welfare loss - even if the value of status is small relative to the value of being in group without congestion.

2.4.2 Status and Social Exclusion

Another question pertains to the 'degree' of participation. Is it ever optimal for a social planner to exclude some agents from joining any group? I refer to this more extreme version of segregation as *social exclusion*.

If agents don't join a group, they are reduced to their stand-alone payoff. This describes a situation where agents don't benefit from any social spillovers; their payoff does not depend on their peers. Since a social planner has to respect incentive compatibility and the utility from group membership net of payments by assumption exceeds the stand-alone payoff, this can only be done through prices. For instance, tuition fees might be high enough to deter some people from acquiring any higher education. They then don't obtain any of the absolute or relative benefits generated by higher education institutions. Because of the restriction to prices, a planner can only exclude agents at the bottom end of the type space - the excluded agents are in an interval $[\underline{w}, \underline{w}_1)$. If some students are deterred from enrolling at university by the tuition fees, then all students with lower ability are also deterred by these prices.

In this section it is shown that while a social planner would not choose to exclude agents if they had preferences over quality[7] or status alone, if both concerns matter, some social exclusion might be a second-best outcome. Social exclusion by a social planner is thus an issue that arises from the interaction of both types of peer effects.

First, a notion of monotonicity of the quality ϕ is introduced that captures the idea that higher types lead to at least weakly higher quality groups - and in this sense ϕ indeed captures the 'quality' of the group. Under weak monotonicity of ϕ, if we observe two social groups such that all types in the support of one group are higher than those in the support of the other, we can conclude that the group with higher types also has (weakly) higher quality.[8] This is, for instance, the case if the quality ϕ is the average type, the lowest or highest type, or a strictly increasing function thereof. The results are robust to variations of this definition that capture similar ideas of monotonicity.

Definition 9 (Weakly monotone ϕ)**.** *The group quality ϕ is weakly monotone if for any two social groups with support over $[\underline{w}_l, \overline{w}_l]$ and $[\underline{w}_h, \overline{w}_h]$ with $\overline{w}_l \leq \underline{w}_h$, the group qualities are such that $\phi_l \leq \phi_h$.*

If ϕ is weakly monotone, then without status concern, excluding agents is always Pareto-dominated:

Lemma 4. *Without status concern (U^q) under weakly monotone quality, for any incentive compatible group provision $m^q(w)$ inducing a partition $\mathscr{I} = \{w_1, ..., \overline{w}\}$ where*

[7]At least under some assumption on ϕ, that make it non-decreasing in agents' types - see Assumption 9.

[8]Given a $\mathscr{L}$, ϕ induces an order of measurable subsets in the agent space $(\Omega, \mathscr{B})$. If weak monotonicity holds, then this ordering is aligned with the ordering of those sets induced by inf and sup in W if inf and sup induce the same order for these sets.

$w_1 > \underline{w}$, there exists an incentive compatible provision $m^{q'}(w)$ such that the equilibrium under $m^{q'}(w)$ Pareto-dominates that under $m(w)$.

Without status concern, there is no benefit in pricing agents out of the market. The social planner can always increase aggregate welfare by offering these agents a separate group or merging them into the existing one. Under monotonicity of ϕ, every partition of W into convex sets can be sustained as an equilibrium.[9] So if one partition Pareto-dominates another, there exists a price vector such that the corresponding social groups can be sustained as an equilibrium in the agents' game. Note that not only has the planner an incentive to offer prices such that all agents participate, this also increases the payoffs of all agents directly as prices are lower under full participation.

When agents have preferences over status only (U^r), the situation is almost opposite: in equilibrium there can at most be one group. Nevertheless, social exclusion is also not optimal because of the inefficient allocation of rank as well as the welfare loss from reducing some agents to their stand-alone payoff.

Result 7.3. *Social exclusion is not optimal when agents have preferences over status only.*

Proof. This follows directly from Corollary 7.2 where it is shown that under status-only preferences, full-participation is the unique welfare-maximizing equilibrium. □

Even when agents have preferences over both quality and status, the first-best outcome where the planner does not have to respect agents' incentive constraints, would also not produce any exclusion. However, since for the constrained problem the planner is restricted to agents truthfully reporting their type and thus self-selecting into groups, exclusion can be a second-best outcome.

Proposition 8. *For any incentive compatible group provision $m(w)$ with partition $\mathscr{I} = \{w_1, ..., \overline{w}\}$ where $w_1 > \underline{w}$, the full-participation group partition $\mathscr{I}' = \mathscr{I} \cup \{\underline{w}\}$ achieves strictly higher aggregate welfare. Under status concern, no provision $m'(w)$ might exist that sustains $\mathscr{I}'$ as an equilibrium.*

We can develop the argument further and show that given some W and some preferences that include status concern, we can always find a probability measure over

[9] To see this, note that for any two adjacent groups g and h with support over $[\underline{w}_g, w^*]$ and $[w^*, \overline{w}_h]$, weak monotonicity guarantees that $\phi_h \geq \phi_g$. Because of supermodularity in w and ϕ, the willingness to pay for quality is strictly increasing in w. Since prices are determined by the utility difference of the cut-off type, we know that without status concern, the increase in utility from the difference $\phi_h - \phi_g$ is higher for w^* than for all $w < w^*$. There exist prices p_g and p_h such that all $w \geq w^*$ prefer membership in h while all $w < w^*$ prefer membership in g.

W and a quality function ϕ, such that at least some social exclusion is a second-best outcome.

Proposition 9. *For every U and W, there exists a weakly monotone quality function ϕ and measure $\mathscr{L}$ such that the exclusion of some agents is a second-best outcome.*

Looking at this from a different angle, let's consider a population where everybody is a member of a social group in the welfare maximizing, stable equilibrium. What if this population is joined by new agents? To keep it simple, suppose we extend the support at the low end of types. Would a social planner offer prices that make these new agents participate? Without status concern and under weak monotonicity of ϕ, the answer is yes. A social planner can always increase aggregate welfare relative to the initial equilibrium by offering prices such that the new agents form a separate group. The new welfare maximizing equilibrium also has full participation. This is not the case with status concern. The social planner might not want the agents to join any of the existing groups and, as the following result demonstrates, there always exists a measure over this extended set such that they cannot be offered a separate group. Even though there might be a larger measure of agents in the population, aggregate surplus beyond the stand-alone payoff does not necessarily increase.

Corollary 9.1. *For every equilibrium group partition $\mathscr{I}$ of W, there exists a measure $\mathscr{L}'$ over $W' \equiv [\underline{w}', \underline{w}] \cup W$ such that $\mathscr{I}' = \mathscr{I} \cup \{w'_0, ...\}$ is not an equilibrium group structure for any $w'_i \in [\underline{w}', \underline{w}]$ and any prices.*

If the support is extended at the top end, a similar situation can arise. Social exclusion occurs because of the effect some agents have on quality. This becomes a stronger force when there are agents that value quality sufficiently highly. So if there is full participation in the welfare maximizing equilibrium for some W and if this W is extended by some interval $(\overline{w}, \overline{w}']$, depending on the preferences and ϕ, this might lead to a positive degree of social exclusion. While here at least for monotone ϕ aggregate surplus necessarily increases - noting that full participation could be maintained - agents with low type can be left reduced to their outside option.

2.5 Privatization

What if the incentives of the group provider are less aligned with the agents' preferences? For instance, an authority might be maximizing revenue instead of welfare. Tuition fees might be set very differently if the objective is to maximize receipts rather than benefits from education. Alternatively, we can think of an authority handing over this task to one or several private companies. What are the consequences of such a privatization? Moving away from aggregate welfare as an objective, the next section focuses on firms supplying groups to maximize profits. As before, a provider posts a set of prices and then agents decide which group to join. As will be shown, if preferences are separable in rank and quality, status concern unambiguously lowers the increase (or magnifies the decrease) in revenue from splitting any given group - segregation is less beneficial. But we reach some negative result for the comparison with the social planner: no clear conclusion can be drawn whether a monopoly leads to more or less segregation or social exclusion. As an example shows, a monopolist might exclude too few agents relative to the constrained welfare maximum. I provide sufficient conditions when this is the case.

Competition, in comparison, leads to lower prices and less social exclusion in some cases but a general statement is not possible without making stronger assumptions on the interactions among firms and consumers. Nevertheless, similar to the previous results, the presence of status concern tends to exacerbate social exclusion.

2.5.1 Monopoly

A natural starting point for this line of inquiry is to examine the monopolist's problem. Just like in the social planner's case, the monopolist offers an incentive compatible, individually rational group provision $m(w) = \big(y(w), p(w)\big)$.[10] Since the provision of groups is assumed to be free, a monopolist offers a group provision that maximizes revenue.

We can state the monopolist's problem as:

$$\begin{aligned} &\max_{y(w),p} \sum_G p_g \mathscr{L}_g(W) \\ &s.t. \quad U(w, m(w), \mathscr{L}_A) \geq U(w, m(w'), \mathscr{L}_A) \quad \forall w, w' \in W \\ &\qquad\;\; U(w, m(w), \mathscr{L}_A) \geq \underline{u} \quad \forall w \in W \end{aligned} \tag{2.11}$$

where $y(w)$ is a function that assigns an agent of type w an action $g \in A$, and $\mathscr{L}_A$ is the vector of social groups generated by this assignment.

[10]As before, incentive compatibility requires that $p(w)$ is constant over the support of each active social group and we can thus focus on prices for each group rather than each agent.

We first look at the problem for preferences that are separable in ϕ and r in the form:

$$U(w, m(w), \mathscr{L}_A) \equiv u\left(w, \phi_{y(w)}\right) + v\left(w, r_{y(w)}(w)\right) - p(w)$$

Differences in group quality do not affect the importance of rank. As discussed before, this fulfils Assumption 4. With this restriction, we can establish that under status concern, segregating any group increases revenue by less than it would under U^q.

Suppose there exist a group provisions $m(w), m^q(w)$ for preferences with and without status concern such that the partition $\mathscr{I}$ represents the associated equilibrium group structure. Suppose further that $\mathscr{I}'$ is a strictly finer partition than $\mathscr{I}$, and there are incentive compatible provisions $m'(w), m^{q\prime}(w)$ that induce $\mathscr{I}'$. Furthermore, prices for each group structure are such that profits are maximized given the group structure.

Proposition 10. *If utility is separable in ϕ and r, then offering $m'(w)$ achieves higher profits than $m(w)$ under status concern only if offering $m^{q\prime}(w)$ achieves higher profits than offering $m^q(w)$ without status concern.*

If group quality does not interact with status directly,[11] the incentive to segregate is weaker for preferences with status-concern; just as in the social planner case. As status concern 'increases' in importance, this effect gets stronger. To demonstrate this, we can parametrize status concern as before:

$$U_\alpha(w, m(w), \mathscr{L}_A) \equiv u\left(w, \phi_{y(w)}, r_{y(w)}(w)\right) + \alpha v\left(w, r_{y(w)}(w)\right) - p(w)$$

And again, let $m_\alpha(w)$ and $m_\alpha{}'(w)$ refer to the incentive compatible, revenue maximizing group provision that, under status concern with parameter α, lead to the partitions $\mathscr{I}$ and $\mathscr{I}'$, where $\mathscr{I}'$ is finer than $\mathscr{I}$:

Corollary 10.1. *The group provision $m_\alpha{}'(w)$ achieves higher profits than $m_\alpha(w)$ under U_α only if $m_{\hat{\alpha}}{}'(w)$ achieves higher profits than $m_{\hat{\alpha}}(w)$ under $U_{\hat{\alpha}}$ for all $\hat{\alpha} \in [0, \alpha]$.*

Note that for this result, we only assumed separability of the component that magnifies status concern. Even when no clear-cut ordering between preferences with and without status concern is possible, we can still order the effects within this class of preferences. Naturally, this includes the case where utility is separable in ϕ and r. Stronger status concern weakens the incentive to segregate for a private provider. Suppose, for example, a private company is tasked with developing a housing project. If

[11] I.e. $\frac{\partial^2}{\partial\phi\partial r} u(w, \phi, r) = 0$

people have strong preferences over their relative standing in their housing community, then the revenue from offering segregated communities is lower than if they only cared about the quality of their community. The stronger the concern for rank relative to quality, the more beneficial it is to offer an inclusive community.

Without separability, the result in Proposition 10 holds for sufficiently weak interaction between ϕ and r. In the general case, the outcome is more ambiguous and depends on the magnitude of the interaction as well as the relative changes in group quality. Intuitively, this is because splitting a group might lead to lower quality for one and higher quality for the other group. Status concern generally lowers price differences and thus the benefit from splitting a group. But since r and ϕ are complements, the reduction in price at the low end of the group might be less under status concern. In other words, the lower quality impacts the agent with rank 0 in the lower quality group less because of his low rank. And since the utility of this agent determines the price of the low-quality group, the drop in revenue caused by the reduction in quality is less. In some cases, this could compensate for the negative revenue effect rank has for the high-quality group. Nevertheless, we can still identify some cases where a clear comparison is possible.

The table below summarizes the different cases that can occur when 'splitting' a group with support over some interval into a high- and a low-type group. In particular, suppose a monopolist offers a modified provision such that a group with support over $[\underline{w}_g, \overline{w}_g]$ and quality ϕ_g is instead separated into two groups with support over $[\underline{w}_l, \overline{w}_l]$ and $[\underline{w}_h, \overline{w}_h]$ where $\underline{w}_l = \underline{w}_g$, $\overline{w}_h = \overline{w}_g$, and $\overline{w}_l = \underline{w}_h$ with qualities ϕ_l and ϕ_h. The effect of this group provision on overall revenue depends on the difference between the new qualities ϕ_l, ϕ_h and the initial group quality ϕ_g. As is summarized below, whether or not we can in this case establish an ordering of the effect on revenue between preferences with and without status concern depends on the sign of the quality differences. Of course, for any such new structure to be an equilibrium, the new provision needs to be incentive compatible which necessarily requires $\phi_l < \phi_h$.

	ϕ_l	ϕ_h	*increase in revenue: status vs. no-status*
(1)	$< \phi_g$	$< \phi_g$	ambiguous
(2)	$< \phi_g$	$> \phi_g$	ambiguous
(3)	$> \phi_g$	$> \phi_g$	lower under status concern
(4)	$> \phi_g$	$< \phi_g$	not incentive compatible

Table 2.1: Comparison of revenue effects - proof see appendix

A clear-cut comparison is only possible if the split leads to an increase in quality for both groups - case (3). Offering the corresponding group provision increases revenue under status concern only if it increases revenue without status concern. Interestingly, even though this split has the most positive welfare effect in terms of quality, status concern unambiguously reduces the return from offering this provision. A monopolistic provider could cause a large welfare loss. But it also highlights that a monopolist might induce less segregation than a social planner.
In cases (1) and (2), a comparison is less straight-forward due to the nature of the complementarities. Observe that a drop in utility for a type $\underline{w}_g$ and an increase in utility for a type $\overline{w}_g$ both (weakly) reduce revenue.[12] Compared to preferences without status concern, the revenue effect of a drop in quality is less negative at $r = 0$. Or in other words, the price difference $p_g - p'_l$ that makes this provision incentive compatible at $\underline{w}_l$ is smaller. A low ranked agent is less sensitive to a drop in quality. At the same time, an increase in quality at $r = 1$ has a more negative negative effect on revenue. High ranked agents are more sensitive to changes in quality. In all cases, the price difference $p'_h - p_g$ is less under status concern. This leads to the ambiguity in the overall effect which depends on the relative magnitudes of all these factors.

Social Exclusion - Monopolist

Concerning social exclusion, the effect of status concern - comparing preferences U and U^q - is twofold. On the one hand, excluding an entire group that could be offered in equilibrium is more beneficial under status concern. However, the increase in revenue from lowering the price at the lowest cut-off - and thus offering a group to more agents - is generally greater when agents' have status concern. Overall, this means the presence of status concern can lead to more or less social exclusion.

The first effect - the greater benefit of excluding an entire group - is captured by the following result: Suppose for preferences with and without status concern there exist incentive compatible group provisions $m(w), m^q(w)$ and $m'(w), m^{q\prime}(w)$ inducing the partitions $\mathscr{I} = \{w_1, ..., \overline{w}\}$ and $\mathscr{I}' \equiv \underline{w} \cup \mathscr{I}$ respectively - i.e. there is less social exclusion under $\mathscr{I}'$. Suppose further the provisions maximize profits given the group structures.

Result 10.2. *Offering $m'(w)$ instead of $m(w)$ leads to higher revenue under status concern, only if offering $m^{q\prime}(w)$ instead of $m^q(w)$ leads to higher revenue without status concern (U^q).*

[12] The former has a strict negative effect while the latter only takes effect for $\overline{w}_h < \overline{w}$

To see this, notice that the intra-group utility difference

$$U(\overline{w}_g, m(\overline{w}_g), \mathscr{L}_A) - U(\underline{w}_g, m(\underline{w}_g), \mathscr{L}_A)$$

is larger under status concern due to the difference in rank. This means by excluding the lowest group, all prices can be raised by that intra-group difference. The logic is the same as in the classic monopoly screening problem of Mussa and Rosen (1978); serving low types has a negative effect on the revenue from higher value types. When agents have status concern, this effect is stronger. The monopolist can raise prices by more if a group is excluded - i.e. if $p(w)$ and $q(w)$ are the price schedules with and without status concern then for all non-excluded w, we have $q(w) - q'(w) < p(w) - p'(w)$.

However, there are cases when a separate group couldn't even be offered in an incentive compatible way. Exclusion is then a question about the effect of marginal changes in the lowest cut-off on revenue. Looking at any given group structure: raising the lowest price to exclude more agents has a less positive effect on revenue when agents care about status since the low rank of a cut-off agent reduces the effect from potential changes in quality (and the effect of the increase in w itself). This is due to the complementarity in rank, type and quality. If a monopoly provider raises the tuition fees of all universities, some students choose not to attend. Given the same change in price, more students are excluded under status-concern since the students' valuation at the cut-off is 'depressed' by the low rank. Status concern can thus, in principle, lead to more or less exclusion if a monopolist serves the market.

The direct comparison between a monopolist and a social planner on the basis of social exclusion in equally ambiguous. But this means a monopolist might set a lower price for the lowest quality group and thus exclude fewer agents than a social planner. This is beneficial to agents with low type. As discussed earlier, the social planner might exclude a set of agents if they have a sufficiently negative influence on the quality of the adjacent group but they cannot be offered a separate, incentive compatible group. The monopolist's incentives, in contrast, are independent of whether or not the agents could be offered a separate group and only depend on the influence on the utility of the cut-off type - as this determines prices. This means the monopolist might exclude fewer agents than the social planner. A numerical example for exactly this case is presented later. The following result summarizes this argument. Note that here separability of status and quality is no longer required.

Suppose $\mathscr{I} = \{w_1,...,\overline{w}\}$ is an equilibrium group structure with $w_1 > \underline{w}$ for some $m(w)$ and $\mathscr{I}' = \{w_1',...\overline{w}\}$ is an equilibrium group structure with $w_1' < w_1$ for some $m'(w)$ and $w_i' = w_i$ for all $i > 1$ - the group structure only differs in the cut-off type of the lowest group; under $\mathscr{I}'$ fewer agents are excluded.

Result 10.3. *If $u(\overline{w},\phi_h,0) - u(\underline{w},\phi_l,0)$ is sufficiently close to 0 for all $\phi_h \geq \phi_l$, and if for any $[\underline{w}_g,\overline{w}_g] \subset W$, the quality ϕ_g is weakly increasing in $\underline{w}_g$, then aggregate welfare is higher under $m'(w)$ than $m(w)$ only if profits are higher under $m'(w)$ than $m(w)$.*

If excluding agents has a very small effect on the utility of the cut-off type, then doing so has little benefit to the monopolist. An increase in price will lead to a decrease in revenue. To the social planner, however, excluding agents might still be beneficial as long as it raises the utility of (some) other members of the group. In that case, a monopolist prefers excluding fewer people than a social planner. Intuitively, if the student with the lowest relative ability in a university benefits very little from interaction with the other students - maybe because he socializes less - then a monopoly provider cannot benefit from making the university more exclusive. In comparison, an authority maximizing welfare might still prefer to do so if it sufficiently benefits students that have a higher relative standing.

As the previous discussion highlights, comparing the optimal provisions of a social planner and a monopolist does not generally allow for straight-forward results if agents care about their status. Regarding the degree of segregation, a monopolist might achieve higher profits with a strictly finer group structure even if this lowers quality in all of the more segregated groups. On the other hand, a monopolist might prefer a coarser provision even though there exists a provision that achieves a finer group structure with higher quality in all of the more segregated groups. Similarly, a monopolist might induce less or no social exclusion while a social planner might choose to exclude agents. The ambiguous interaction between type, quality and rank denies any more clear-cut comparison without stronger assumptions on preferences.

2.5.2 Competition

Instead of a monopolist, the market for social groups might be served by several competing firms - municipalities competing for inhabitants or universities competing for students with the goal to maximize revenue. This section provides a brief discussion of some of the implications of such a competitive environment. It is shown that at least for full-participation outcomes, competition leads to lower prices and potentially even subsidized memberships. Furthermore - unlike in a monopoly setting - no set of agents is excluded if they could be offered a separate group. Competition thus reduces social exclusion in some cases. Nevertheless, status concern allows for equilibria with social exclusion even though without status concern, offering the excluded agents a separate group would be a Pareto-improvement. Competition does, however, not generally lead to more or less segregation or social exclusion. Ultimately, stronger assumptions on the agents' coordination would be necessary to derive sharper predictions.

Suppose instead of a single firm, there is a large (countable) number of potential providers. As in the monopoly case, every firm can offer a menu of prices and assignments. We combine all these offers into the set of group provisions:

$$\{m_n(w)\}_{n\in N}$$

where $m_n(w)$ is the offered group provision of firm n. Since each firm offers their own, 'unique' groups, agents joining group $y_i(w)$ according to some provision $m_i(w)$ are members of a different group than agents joining group $y_j(w)$ according to some $m_j(w)$ with $i \neq j$. The timing is unchanged: the firms simultaneously offer group provisions and maximize revenue given the other firms' and the agents' strategies. Then the agents simultaneously choose a supplier and report their type. This is still referred to as the agents' game. Different agents can, of course, go to different suppliers so that, for example, some set $W_1 \subset W$ is supplied by firm i while a different set $W_2 \subset W$ is supplied by firm j. Agents of the same type, however, are restricted to choose the same supplier. This ensures that all equilibria in this game are stable in the sense of Chapter 1.

To keep this closely aligned to the monopoly setting, this is modelled as a mechanism μ that, given the firms' offers and the agents' report of their type, recommends a provision m_n. This μ is called the market provision:

Definition 10. *The market provision μ is a $\mathscr{B}_W$-measurable function that assigns every type $w \in W$ a group and price consistent with some group provision $m_{n(w)}(w) \in \{m_n(w)\}_{j\in N}$:*[13]

$$\mu(w) = m_{n(w)}(w)$$

[13]This reduces to the monopoly case for $N = 1$ and μ being the identity function $\mu(w) = m(w)$.

Given the continuum of players, the coordination aspect needs to be carefully considered as single-player deviations on the agents' side have no bearing on the aggregate outcome. For example, if all agents decided to join a particular group, there cannot be a beneficial deviation to another group for any single agent. Without any restriction, any incentive compatible group structure can be an equilibrium outcome because of the coordination aspect in the agents' game. In the previous setting, we circumvented the problem by selecting the seller-optimal equilibria through a mechanism. With a large number of competing firms, a focus on seller-optimal outcomes seems inappropriate. It appears natural to shift some of the advantage towards the agents. We therefore focus on equilibria that are Pareto-efficient from the agents' perspective:

Definition 11. *A market provision μ given a set of provisions $\{m_n(w)\}_{n\in N}$ is Pareto efficient in the agents' game if there is no other market provision μ' given $\{m_n(w)\}_{n\in N}$ such that all agents are weakly better-off and a positive measure of agents is strictly better-off.*

The analysis focuses on all market provisions that are incentive compatible, individually rational, and Pareto-efficient in the agents' game and that are a Nash equilibrium from the firms' perspective. Formally, these are the subset of subgame-perfect equilibria in which the market provision constitutes a Pareto-efficient correlated equilibrium in the agents' game.[14] As again incentive compatibility requires that active prices[15] are constant over the support of a social group, the equilibrium outcomes can be described in terms of the vector of active social groups $\mathscr{L}_A$, the partition $\mathscr{I}$ generated by the social groups, as well as the price vector p - with the maintained convention that $p_i \geq p_j$ if $i > j$. In this context, we can distinguish between settings where prices are restricted to be non-negative $p(w) \in \mathbb{R}^+$ and settings where subsidies for some groups are feasible $p(w) \in \mathbb{R}$.[16]

The focus on Pareto-efficiency in the agents' game allows us to derive a Bertrand-type result. In any equilibrium where no agents are excluded, the lowest (active) price has to be 0. If not, some other firm could offer a price vector that undercuts the other firm while maintaining the same group structure. This would necessarily be a Pareto improvement in the agents' game; the market provision would have to recommend

[14]It is not necessary that each firm's provision is incentive compatible as only a subset might choose this firm in equilibrium. Thus only the market provision needs to be incentive compatible. There is, however, no loss in focusing only on the prices and groups chosen by a measurable set of agents.

[15]Recall that 'active' prices refers to those prices that are associated with a group that is chosen by a measurable set of agents.

[16]In the case of a monopolist supplier or social planner, a restriction to non-negative prices is inconsequential as all incentive compatible group structures can be implemented with non-negative prices and subsidies are never profit maximizing for a monopolist.

this firm's provision to all agents and this would allow that firm to capture the entire market.

Result 10.4. *In a competitive setting with $p(w) \in \mathbb{R}^+$, any full-participation equilibrium that is Pareto-efficient in the agents' game has the lowest active price $p_1 = 0$.*

If the restriction to non-negative prices is removed, it becomes possible to subsidize membership in lower quality groups. This ensures that profits are reduced to 0. Naturally, this is only possible in a full-participation outcome; with $p_1 \leq 0$, no agents can be excluded. Furthermore, it implies that all full-participation equilibria are payoff-equivalent to one where a single supplier serves the entire market.[17] The intuition is as follows: At least the highest quality group has to be offered at a positive price to ensure non-negative profits. If one supplier offers all groups that are priced positively, then this supplier can also provide the highest membership subsidies. The equilibrium that is Pareto efficient from the agents' perspective is the one with the lowest prices - this is the case for a single supplier. In general, the number of firms providing active groups in a stable equilibrium with full-participation cannot be larger than the number of active (weakly) positive and (weakly) negative prices. If, for example, there two active groups, there can be only one supplier. However, this does not imply that there will generally be full participation. In fact, allowing for non-negative prices has no influence on equilibria with a positive level of social exclusion.

Result 10.5. *In a competitive setting with $p(w) \in \mathbb{R}$, any full-participation equilibrium Pareto efficient in the agents' game with more than one active social group has $p_1 < 0$. The number of firms providing active groups is bounded by both the number of non-positive active prices and the number of non-negative active prices. All firms make 0 profits.*

Competition also rules out some cases of social exclusion. In particular, if a set of agents at the low end of W can be offered a separate group, then excluding them cannot be an equilibrium. A firm could offer a group provision that leads to a Pareto improvement for all agents. This differs from the monopoly setting where a monopolist might have an incentive to not offer this group in order to increase revenue. Nevertheless, as discussed before, if status concern is sufficiently strong, it might not be possible to offer such a separate group that maintains incentive compatibility for all other social groups.

[17] To maintain equilibrium, at least one other firm has to offer the same provision but it is not recommended my the market provision to a measurable set of agents.

Result 10.6. *Suppose there is an incentive compatible market provision μ inducing group structure $\mathscr{I} = \{w_1, ..., \overline{w}\}$ and μ' inducing group structure $\mathscr{I}' = \{w_0, w_1, ..., \overline{w}\}$. Under competition, $\mathscr{I} = \{w_1, ..., \overline{w}\}$ is not an equilibrium group structure.*

Competition is, in some sense, good news to the agents in at least the full participation equilibria. It necessarily drives down all group prices relative to the corresponding monopoly prices. On top of that, it also rules out some equilibria where agents are excluded: under competition, no set of agents is excluded if they could at least be offered a separate group without changing the remaining group structure. Nevertheless, despite the focus on Pareto efficient equilibria, competition does not necessarily lead to an increase in aggregate welfare or even just a reduction in social exclusion. In particular, in settings where many equilibria are on the agents' Pareto frontier, there can be equilibria with higher prices and lower welfare compared to the monopoly outcome. The example in the following section details such a case.

The role of status here is more subtle. Without status concern, the number of active social groups could approach infinity. This can pose a problem for existence. As was shown in Proposition 4, the number of active social groups under status concern in any equilibrium is bounded. In this sense, competition does not immediately lead to a problem with existence. Furthermore, the presence of status concern implies that changes in social groups have more heterogeneous effects on agents. Shifting the group boundary or splitting a group into several groups can increase the utility of some agents and reduce that of others due to the accompanying changes in rank. As a consequence, many different equilibria can lie on the agents' Pareto frontier.

In comparison to a monopoly setting, no social exclusion occurs in cases where agents can be offered a separate group. But if status concern is sufficiently strong, this might not be incentive compatible. Without status concern, such a group would always be offered in a competitive market. Status concern exacerbates social exclusion even in a competitive setting. In fact, there can be competitive equilibria with more agents excluded than both in the planner and the monopoly environment. The overall effect of competition on welfare is thus ambiguous. It is dependent on how exactly the agents (can) coordinate and further research would be desirable to examine the competitive provision in more detail.

2.6 Numerical Example

The following example illustrates previous claims about the social exclusion induced by a social planner and a monopolist: it might be a second best to exclude some agents from participating and, in fact, more agents than in a monopoly setting. Furthermore, it is demonstrated that even in a competitive environment, agents might be excluded - and indeed potentially more so than under the two other market structures.

In the example below, types are distributed over $[1,2]$ according to a (truncated) Pareto distribution with parameter 1. The example is set-up such that there can be at most one group in equilibrium. Depending on the relative importance of quality and rank, we can create scenarios in which both a social planner and a monopolist exclude some agents, only the planner excludes agents or neither excludes any agents. In one case, the monopolist offers a price that all agents accept while the social planner offers a price such that all agents with type lower than $\underline{w}_1 \approx 1.22$ prefer their outside option. In each of the cases, the quality of the group is determined by the lowest type in the group. Accordingly, if we compare this to a setting where agents have preferences over status only (e.g. $r = 1/2$ for all w), segregating a given partition further always leads to an increase in welfare and there always exists a price vector that makes this incentive compatible.

Example 1: Types are distributed according to a truncated Pareto distribution (shape $a = 1$) over $[1,2]$. Preferences are represented by:

$$U(w, g, \mathscr{L}_A) = (\alpha_1 + \beta_1 w^{\gamma_1})\phi_g + (\alpha_2 + \beta_2 w^{\gamma_2})\ln(r_g(w) + \frac{1}{2}) - p_g \tag{2.12}$$

The quality of a social group $\mathscr{L}_g > 0$ is given by:

$$\phi_g = \frac{1}{2} + \frac{1}{2}\underline{w}_g \tag{2.13}$$

Table 2.2: Optimal levels of social exclusion

(1)	$\gamma_1 = 15.48$	$\gamma_2 = 15.50$	eq. group structure
social planner	$\underline{w}_1^{SP} \approx 1.22$		$\mathscr{I}^{SP} = \{1.22, 2\}$
monopolist	$\underline{w}_1^{M} = 1$		$\mathscr{I}^{M} = \{1, 2\}$
(2)	$\gamma_1 = 15.00$	$\gamma_2 = 15.00$	
social planner	$\underline{w}_1^{SP} = 1$		$\mathscr{I}^{SP} = \{1, 2\}$
monopolist	$\underline{w}_1^{M} = 1$		$\mathscr{I}^{M} = \{1, 2\}$
(3)	$\gamma_1 = 16.00$	$\gamma_2 = 15.00$	
social planner	$\underline{w}_1^{SP} \approx 1.44$		$\mathscr{I}^{SP} = \{1.44, 2\}$
monopolist	$\underline{w}_1^{M} \approx 1.88$		$\mathscr{I}^{M} = \{1.88, 2\}$

For all cases: $\alpha_1 = 560$, $\alpha_2 = 15\ln(2)^{-1}$, $\beta_1 = 1$ $\beta_2 = \ln(2)^{-1}$

Example 1 [Table 2.2] gives an indication of the underlying logic for the exclusion of agents. Social exclusion occurs in cases (1) and (3); in (1) only the social planner excludes agents and in (3), both the planner and the monopolist exclude agents. In case (1), the complementarity between type and quality is stronger than in case (2) where no exclusion is optimal. In principle, this provides an incentive to exclude some agents in order to raise the quality, given that the quality is determined by the lowest type. At the same time, the complementarity between rank and type is also stronger than in (2). Excluding agents lowers the rank of almost everyone in the group. This is an offsetting force to the increase in quality. While on aggregate the quality dominates - hence the positive social exclusion in the planner case - the monopolist does not take into account the higher utility of agents strictly inside the support. In this case, an increase in price does not sufficiently increase revenue generated by the agents joining the group to compensate for the loss of revenue from the excluded agents. A monopolist prefers to serve all agents as in (2). In case (3), the complementarity is strong enough that the incentive to raise quality by excluding agents dominates for both the planner and the monopolist.

Extending the example to a competitive environment, we first note that since in any equilibrium there can be at most one group, competition does not necessarily lead to 0 profits. However, there does exist a full-participation equilibrium where $p_1 = 0$ with a single group - which follows from Result 10.4. But as indicated, there are other equilibria on the Pareto frontier where there is social exclusion. In fact, there is a continuum of equilibria with $\underline{w}_1 \in [1, 1.76]$. At high levels of exclusion, welfare is even lower than in the monopoly case.

Table 2.3: Social exclusion in a competitive environment

(1*)	$\gamma_1 = 15.48$	$\gamma_2 = 15.50$	eq. group structure
competition	*any* $\underline{w}_1^C \in [1, 1.76]$ *is PE in the agents' subgame*		$\mathscr{I}^C = \{\underline{w}_1^C, 2\}$

As before: $\alpha_1 = 560$, $\alpha_2 = 15\ln(2)^{-1}$, $\beta_1 = 1$ $\beta_2 = \ln(2)^{-1}$

2.7 Conclusion

In this model, a large number of agents observe a set of membership prices and simultaneously decide which group to join. Group membership gives them access to social spillovers generated by the other members. In particular, agents have preferences over the composition of the group (the quality) and their rank in the distribution of types within the group (their status). The novel aspect of this paper is the exploration of the role of status concern in such a setting. The focus lies on the consequences for segregation and social exclusion.

This chapter qualifies the following two broad claim: Firstly, status concern reduces the benefit in terms of welfare and revenue from segregating agents. The incentive to segregate any given group is, in many cases, lower under status concern. This applies - albeit to different degrees - whether the market is served by a social planner, a monopolist or by competitive firms. Secondly, status concern in combination with preferences over the quality of a group can aggravate social exclusion. It can be a second-best outcome to exclude a positive measure of agents. Status concern also creates a stronger incentive for a monopolist to exclude a set of agents even though they could be offered a separate group. And even in a competitive setting, there can be equilibria with a positive degree of social exclusion. Most importantly, by comparing this to the setting in Board (2009), where agents do not have preferences over their rank, we can establish that it is indeed status concern that is driving these results.

Since status concern can lead to different conclusions about the welfare and revenue maximizing group structures, determining whether agents have preferences over status in a given setting can be crucial for empirical investigations. If we observe that part of the population is excluded from a certain type of social group (e.g. tuition fees being set such as to exclude some potential students), we cannot immediately conclude that authority is maximizing profits as opposed to welfare. It could be a constrained welfare optimum to do so. If the quality of groups is monotone in types - i.e. the higher the types the higher the group quality - and there are complementarities between type and quality, then without status concern prices are sufficient to separate agents into any partition of groups. Sorting can be arbitrarily fine. With status concern, this is not feasible. There is a role for screening methods other than prices. Universities, for example, rarely rely on prices alone for admission decisions.

As a final comment, it is the absence of status concern that allows for a high degree of segregation. An authority might want to actively discourage positional concerns and status thinking. But in this setting, status concern limits the degree to which agents can be separated. Discouraging positional concern might lead to more rather than less segregation - and thus to more unequal access to the semi-public good that is offered through the group. On the other hand, positional concerns can lead to the social exclusion of agents. A policy maker might have to consider this trade-off.

2.8 Appendix

2.8.1 A: Omitted Proofs

Proof of Lemma 3 :

Proof. Suppose not and a measurable set of agents pay different prices for the same group membership. This implies that there are sets $W_1, W_2 \subset W$ with $y(w') = y(w'')$ for all $w', w'' \in W_1 \cup W_2$, and $p(w_1) = p(w')$ for all $w_1, w' \in W_1$ as well as $p(w_2) = p(w'')$ for all $w_1, w' \in W_1$. However, $p(w_1) \neq p(w_2)$ for all $w_1 \in W_1$ and $w_2 \in W_2$. Suppose w.l.o.g. that $p(w_1) < p(w_2)$. This group provision cannot be incentive compatible as agents in W_2 could obtain strictly higher utility by reporting a type in W_1 as

$$u(w_2, \phi_{y(w_1)}, r_{y(w_1)}(w_2)) - p(w_1) > u(w_2, \phi_{y(w_2)}, r_{y(w_2)}(w_2)) - p(w_2)$$

since $y(w_1) = y(w_2)$ and misreporting the type does not affect the rank of an agent conditional on being assigned to the same group. The result follows. □

Proof of Proposition 6:

Proof. It follows directly from Lemma 3 that $p(w)$ needs to be constant over the support of each social group.

As $m(w)$ assigns each type $w \in W$ a choice in A, the support of each social group generated by m intersect on a measure 0 set by construction. It follows that if the assignment $y(w)$ is an equilibrium given p, it is a stable equilibrium. Following the equilibrium definition (Definition 2) in Chapter 1, incentive compatibility and individual rationality directly imply an equilibrium in the agents' game. Every incentive compatible and individually rational $m(w)$ is a stable equilibrium in the agents' game.

To show the other direction, note that in an equilibrium in the agents' game, prices are constant for each group by construction since p contains a price for each $g \in A$ (but no more). The equilibrium definition implies directly incentive compatibility and individual rationality. The result follows. □

Proof of Proposition 7:

Proof. I prove the contrapositive: if splitting a group is not beneficial without status concern, it cannot be beneficial under status concern since the benefit from doing is is necessarily less.

Since $\mathscr{I}'$ is a finer group structure than $\mathscr{I}$, there exists at least one active social group $\mathscr{L}_g$ under $\mathscr{I}$ whose support is split into several groups under $\mathscr{I}'$. Suppose the support of g is split into $\mathscr{L}_l$ and $\mathscr{L}_h$ with qualities ϕ_l and ϕ_h. WLOG, assume that $\phi_l < \phi_h$ noting that they can't be equal if m' is incentive compatible. If a refinement does not increase welfare without status concern, then there are $\overline{w}_l = \underline{w}_h \in \mathscr{I}'$ such that:

$$\begin{aligned}\int_{\underline{w}_l}^{\overline{w}_l} u\left(w, \phi_l, r_0\right) dF(w) + \int_{\underline{w}_h}^{\overline{w}_h} u\left(w, \phi_h, r_0\right) dF(w) \\ \leq \int_{\underline{w}_l}^{\overline{w}_h} u\left(w, \phi_g, r_0\right) dF(w)\end{aligned} \tag{2.14}$$

where r_0 is the fixed reference rank of U^q.

Comparing this to preferences with status concern, we note first that it follows from the quality-status-neutrality assumption that for $i \in \{l, h\}$ and for any ϕ:

$$\frac{\partial}{\partial \phi} \int_{\underline{w}_i}^{\overline{w}_i} u\left(w, \phi, r_i(w)\right) dF(w) = \frac{\partial}{\partial \phi} \int_{\underline{w}_i}^{\overline{w}_i} u\left(w, \phi, r_0\right) dF(w)$$

We can conclude that, fixing the ranks agents obtain after the split:

$$\begin{aligned}&\int_{\underline{w}_l}^{\overline{w}_l} u\left(w, \phi_l, r_l(w)\right) dF(w) + \int_{\underline{w}_h}^{\overline{w}_h} u\left(w, \phi_h, r_h(w)\right) dF(w) \\ \leq &\int_{\underline{w}_l}^{\overline{w}_l} u\left(w, \phi_g, r_l(w)\right) dF(w) + \int_{\underline{w}_h}^{\overline{w}_h} u\left(w, \phi_g, r_h(w)\right) dF(w)\end{aligned} \tag{2.15}$$

The final step is to show that the rank re-allocation creates an additional welfare loss which would imply that the right-hand-side of inequality (2.15) is strictly less than the utility in the coarser group g which is:

$$\int_{\underline{w}_l}^{\overline{w}_l} u\left(w, \phi_g, r_g(w)\right) dF(w) + \int_{\underline{w}_h}^{\overline{w}_h} u\left(w, \phi_g, r_g(w)\right) dF(w) \tag{2.16}$$

I show this using an intermediate step with an auxiliary construction that allows to demonstrate that any re-allocation of rank in the manner of the split induces a welfare loss. This is then used to show that splitting any group is less beneficial under status concern - the result follows.

Take a hypothetical group structure with two types of agents, w^- and w^+ where $w^- \leq w^+$. The measure of agents is as in $[\underline{w}_l, \overline{w}_h]$. Denote this $k_0 \equiv \mathscr{L}_W([\underline{w}_l, \overline{w}_h])$. First, we can look at the limit case where $w^- = w^+ = w$ but all ranks are still allocated. Making use of the probability integral transform, we can establish that the distribution of ranks in any group is uniform. The hypothetical aggregate welfare from such a group with quality ϕ can then be written as:

$$k_0 \int_0^1 u(w, \phi, x) dx$$

Since k_0 only acts as a scaling parameter, it can be normalized to 1. The integral can be written in two alternative but equivalent ways where again $w^- = w^+ = w$.
Firstly, for some $x \in (0,1)$:

$$\int_0^1 u(w, \phi, r) dr = \int_0^x u(w^-, \phi, r) dr + \int_x^1 u(w^+, \phi, r) dr$$

and secondly:

$$\int_0^1 u(w, \phi, r) dr = x \int_0^1 u(w^-, \phi, r) dr + (1-x) \int_0^1 u(w^+, \phi, r) dr$$

Now consider a $w^+ > w^- = w$. From the complementarity in w and r, we can conclude that

$$\int_0^x u(w^-, \phi, r) dr + \int_x^1 u(w^+, \phi, r) dr > x \int_0^1 u(w^-, \phi, r) dr + (1-x) \int_0^1 u(w^+, \phi, r) dr$$

as

$$\int_x^1 u_w(w^+, \phi, r) dr > (1-x) \int_0^1 u_w(w^+, \phi, r) dr$$

If there is a difference in types between one group and the other, aggregate utility is higher when ranks a higher in the high-type group.

Next note that for $w^+ \to w$, $u(w^+, \phi_g, x)$ is the lower bound for almost all w with $r(w) \geq x$ in the original group. We can re-write the sums in terms of integrals over types instead of ranks.
Firstly,

$$k_0 \int_x^1 u(w^+, \phi, r) dr = \int_{w^+}^{\overline{w}_h} u(w^+, \phi, r(w)) dF(w)$$

where $r(w) \equiv \dfrac{F(w) - F(\underline{w}_l)}{F(\overline{w}_h) - F(\underline{w}_l)}$.
and secondly:

$$k_0(1-x) \int_0^1 u(w^+, \phi, r) dr = \int_{w^+}^{\overline{w}_h} u(w^+, \phi, r_h^*(w)) dF(w)$$

where $r_h^*(w) \equiv \dfrac{F(w)-F(\underline{w}_h)}{F(\overline{w}_h)-F(\underline{w}_h)}$.

Inspecting $r_h^*(w)$ and $r(w)$, we can conclude that for almost all $w \in [w^+, \overline{w}_h]$, $r_h^*(w) < r(w)$. For $w^+ \to \underline{w}_h$, we can write the difference between the hypothetical utility and the actual sum of utilities as:

$$\Delta_{w^+} \equiv \int_{w^+}^{\overline{w}_h} \big(u(w,\phi,r_h(w)) - u(w^+,\phi,r(w))\big)\, dF(w)$$

and

$$\Delta_{w^+}^* \equiv \int_{w^+}^{\overline{w}_h} \big(u(w,\phi,r_h^*(w)) - u(w^+,\phi,r_h^*(w))\big)\, dF(w)$$

Observe that for almost every w

$$u(w,\phi,r(w)) - u(w^+,\phi,r(w)) > u(w,\phi,r_h^*(w)) - u(w^+,\phi,r_h^*(w))$$

due to the complementarity in r and w and therefore:

$$\Delta_{w^+} > \Delta_{w^+}^*$$

An equivalent argument can be constructed for w^-. Re-writing both summations we get:

$$k_0 \int_0^x u(w^-,\phi,r)\,dr = \int_{\underline{w}_l}^{w^-} u(w^-,\phi,r(w))\,dF(w)$$

and

$$k_0 x \int_0^1 u(w^-,\phi,r)\,dr = \int_{\underline{w}_l}^{w^-} u(w^-,\phi,r_l^*(w))\,dF(w)$$

with $r_l^*(w) \equiv \dfrac{F(w)-F(\underline{w}_l)}{F(\overline{w}_l)-F(\underline{w}_l)}$.

We can again conclude by inspection that for almost all $w \in [\underline{w}_l, w^-]$, $r_h^*(w) > r(w)$. As before, for $w^- \to \underline{w}_h$ we re-write the difference between this hypothetical utility and the actual utility as:

$$\Delta_{w^-} \equiv \int_{\underline{w}_l}^{w^-} \big(u(w^-,\phi,r(w)) - u(w,\phi,r(w))\big)\, dF(w)$$

and

$$\Delta_{w^-}^* \equiv \int_{\underline{w}_l}^{w^-} \big(u(w^-,\phi,r_l^*(w)) - u(w,\phi,r_l^*(w))\big)\, dF(w)$$

Again, due to the complementarity in w and r, we can conclude that $\Delta^*_{w^-} > \Delta_{w^-}$. It follows that

$$\begin{aligned}\int_{\underline{w}_l}^{\overline{w}_h} u(w,\phi,r(w))dF(w) &= k_0 \int_0^1 u(w,\phi,r)dr + \Delta_{w^+} - \Delta_{w^-} \\ &> k_0 \int_0^1 u(w,\phi,r)dr + \Delta^*_{w^+} - \Delta^*_{w^-} \\ &= \int_{\underline{w}_l}^{\overline{w}_l} u(w,\phi,r_l^*(w))dF(w) + \int_{\underline{w}_h}^{\overline{w}_h} u(w,\phi,r_h^*(w))dF(w)\end{aligned}$$

where $\underline{w}_h = \overline{w}_l$.

The alternative assignment of ranks strictly lowers utility. As the effect of a change in ϕ is the same with and without status concern and the ranks necessarily change in the above defined fashion caused by any split in the support of a group, any finer provision has a less positive effect on welfare under status concern than without status concern. As any finer provision can be regarded as an iteration of binary splits, the result follows. □

Proof of Corollary 7.1:

Proof. The result follows from the proof of Proposition 7 if we re-define the Δ_w and Δ^*_w differences accordingly. Δ_{w^+} and $\Delta^*_{w^+}$ are then:

$$\Delta_{w^+} \equiv \int_{w^+}^{\overline{w}_h} \big(u(w,\phi,r_h(w)) - u(w^+,\phi,r(w)) + \alpha(v(w,r_h(w) - v(w^+,r_h(w))\big)\,dF(w)$$

$$\Delta^*_{w^+} \equiv \int_{w^+}^{\overline{w}_h} \big(u(w,\phi,r_h^*(w)) - u(w^+,\phi,r_h^*(w)) + \alpha(v(w,r_h^*(w) - v(w^+,r_h^*(w))\big)\,dF(w)$$

And we observe that $\Delta_{w^+} - \Delta^*_{w^+}$ is, for the same reason as before, positive and increasing in α. The equivalent statement is true for $-(\Delta_{w^-} - \Delta^*_{w^-})$. The result follows. □

Proof of Corollary 7.2:

Proof. It follows from the proof of Proposition 7 that changing the rank allocation for every type w to a lower rank reduces aggregate welfare.
In particular, $\int_{w^-}^{\overline{w}} u(w,r_{w^-}(w))dF(w)$ decreases for a FOSD shift in $r(w)$ which corresponds to an increase in w^- noting that

$$r_{w^-}(w) \equiv \frac{F(w) - F(w^-)}{F(\overline{w}) - F(w^-)}$$

An upward shift in the lower boundary of any group reduces welfare for that group. Furthermore, it follows equally that splitting the support of any group, say the full participation group, strictly decreases welfare. The welfare maximizing equilibrium is the full participation equilibrium. Suppose $m(w) = (g, 0)$ for some $g \in G$. All agents joining group g satisfies the participation constraint as $p(w) = 0$ and since there is no other active group, incentive compatibility holds. This is an equilibrium group provision. □

Proof of Lemma 4:

Proof. Suppose the social planner offers an incentive compatible group provision inducing prices p and partition $\mathscr{I} = \{w_1, ..., \overline{w}\}$ with $w_1 > \underline{w}$.
This implies that the lowest price in p, call this p_1, is greater than 0; otherwise agents with type $w \in [\underline{w}, w_1)$ have a strict incentive to join that group.
Because of the complementarity in type and quality, there exists a $m^{q'}(w)$ with prices p' that maintains partition $\mathscr{I}' = \{w_0, w_1, ..., \overline{w}\}$ as an equilibrium as long as $\phi_0 < \phi_1$, where ϕ_i is the quality of the group with support over $[w_i, w_{i+1}]$. This is guaranteed by weak monotonicity.
Since the quality of all social groups under m^q remain unchanged, aggregate utility from agents with types in $[w_1, \overline{w}]$ does not change. However, agents in $(w_0, w_1]$ receive payoffs strictly above their outside option which raises aggregate utility. This also implies prices of all active groups are now lower than before as $p'_1 \leq u(w_1, \phi_1) - u(w_1, \phi_0) + u(w_0, \phi_0) - \underline{u} < u(w_1, \phi_1) - \underline{u} = p_1$. where $\underline{u}$ is the agents' stand-alone payoff. It must be that $p'(w) < p(w)$ for all $w \in [w_1, \overline{w}]$. All agents that were member of a group under m^q pay a lower price under $m^{q'}(w)$ for the same group membership and all agents in $(w_0, w_1]$ are strictly better-off. This is a Pareto-improvement. □

Proof of Proposition 8:

Proof. It follows almost immediately from Lemma 4 that extending the interval partition at the tail end by adding another element strictly increases aggregate welfare. As $u(w, \phi, r) \geq \underline{u}$ for any $w \in [\underline{w}, \overline{w}]$, $\phi \in [\underline{\phi}, \overline{\phi}]$ and $r \in [0, 1]$, the same applies here. If a p' exists such that $\mathscr{I}'$ is incentive compatible for all w, then this is an equilibrium with higher aggregate welfare and strictly higher utility for almost all agents.
However, such a p' might not exist as is shown in the example in Section 2.6. □

Proof of Proposition 9:

Proof. We construct a ϕ and $\mathscr{L}$ to validate the claim.

Take some type $w^* \in W$ with $w^* > \underline{w}$ and let S be any measurable set in $\mathscr{B}$ and W^S the smallest convex set in W that contains all types w in S.

$$\phi(S) = \begin{cases} \inf W^S + \sup W^S & \text{for} \sup W^S \leq w^* \\ \inf W^S + w^* & \text{for} \sup W^S > w^* \end{cases}$$

This quality function is weakly monotone.
We can now establish that for for any group structure $\mathscr{I} = \{w_1, w_2, ...\}$ with $w_1 \leq w^* < w_2$, there exists no set of prices such that $\mathscr{I}' = \{w_0\} \cup \mathscr{I}$ for any $w_0 \in W$ with $w_0 < w_1$ is a stable equilibrium group structure if w^* is sufficiently close to $\underline{w}$. This implies that there is no incentive compatible group provision $m'(w)$ that achieves the partition $\mathscr{I}'$. This follows from the fact that the quality difference between any two adjacent groups whose quality is weakly less than that of $[w_1, w_2]$ is bounded by the quality difference between the group with support $[w_0, w_1]$ and the group with support $[w_1, w_2]$. This difference is $w^* - \underline{w}$. As $w^* \to \underline{w}$, this difference goes to 0 while the difference $u(w, \phi, 1) - u(w, \phi, 0)$ is bounded away from 0 for any w and ϕ. This means if there is any active group with w^* in the support, there can be no other active group with types lower than w^* in any equilibrium - if w^* is sufficiently close to $\underline{w}$.

Let now $w_1 = w^*$ and the difference $w^* - \underline{w} = \delta > 0$ such that no separate lower quality group can be offered. We can set $\mathscr{L}([w^*, w)$ sufficiently close to 1 for every $w > w^*$ such that excluding w^* is never optimal. We can conclude that any group structure that is an equilibrium in the agent's game can be either written as $\mathscr{I}$ (i.e. no full participation with the lowest-quality group containing w^* in it's support) or as the full-participation structure $\underline{w} \cup \{\mathscr{I} \setminus \{w_1\}\}$. It remains to be shown that offering a group structure of type $\mathscr{I}$ is optimal for some $\mathscr{L}$.

To this end, we first note that ϕ does not depend on $\mathscr{L}$ beyond what is already determined by W. Suppose now $\mathscr{L}([\underline{w}, w^*]) = \epsilon$. Offering $\mathscr{I}$ instead of $\mathscr{I}' = \{\underline{w}\} \cup \{\mathscr{I} \setminus \{w^*\}\}$ lowers the rank of almost every agent with type $[w^*, w_2]$ but increases the quality of that group by $w^* - w_0 = \delta > 0$. While the quality change is independent of $\mathscr{L}$, the rank decrease for every agent is at most ϵ. Finally, we note that $w_2 - w^*$ cannot be arbitrarily small because if $w_2 \to w^*$ then the quality difference between group g_2 and g_1 goes to 0 which cannot be incentive compatible.

Suppose $w_2 < \overline{w}$ meaning that there is a group g_2. In this case, let

$$\underline{w}_2 \equiv \inf\{w \in W : u(w, \phi_2(w), 0) - u(w, \phi_1, 1) > 0\}$$

where $\phi_2(w)$ is the quality of the group with support over $[w_2, w_3]$ with $w_2, w_3 \in \mathscr{I}$. This is the lower bound such that $u(w_2, \phi_2, 0) - u(w_2, \phi_1, 1) \geq 0$ for all $w_2 \geq \underline{w}_2$. Note that this must be strictly greater than w^* as u is strictly increasing in rank and $\phi_2 \to \phi_1$ for $w_2 \to w^*$. If there is no g_2 meaning that $\mathscr{I}$ only describes one group, set $\underline{w}_2 = \overline{w}$.

Due to the continuity of U, we can conclude that there exists a small enough $\epsilon > 0$ such that for every equilibrium group structure $\mathscr{I}$ with $w^* = \min \mathscr{I}$, the increase in aggregate utility for agents in $[w^*, w_2]$ outweighs the loss from excluding $[\underline{w}, w^*)$ and the loss in rank of agents in $[w^*, w_2]$ given that both the loss in aggregate utility as well as the loss in rank are decreasing in ϵ while the increase in quality remains the same and $\mathscr{L}([w^*, w_2]) \geq \mathscr{L}([w^*, \underline{w}_2]) > 0$ for any equilibrium partition $\mathscr{I}$. We conclude that for every equilibrium $\mathscr{I}$, offering the full-participation group structure $\mathscr{I}' = \{\underline{w}\} \cup \{\mathscr{I} \setminus \{w^*\}\}$ strictly lowers aggregate welfare under the proposed quality function. The result follows. □

Proof of Corollary 9.1:

Proof. Let $\mathscr{I}$ be a group partition of W for an equilibrium $m(w)$ with a corresponding price vector p. Suppose $\underline{w} \in \mathscr{I}$, meaning there is full participation. This is without loss since otherwise we can adjust the definition of W accordingly. Let's extend W at the bottom-end by an interval $[w', \underline{w})$ and call $[w', \underline{w}) \cup W \equiv W'$. Let $\mathscr{L}'$ be the measure over W' with the obvious restriction that $\mathscr{L}'(A) = \mathscr{L}(A)$ for all $A \subseteq W$.

Consider a partition $\mathscr{I}' = \{w'\} \cup \mathscr{I}$. As shown before, this strictly increases welfare and can thus be implemented by the social planner if there exists a p' such that the corresponding group assignments are incentive compatible.

Let $\mathscr{L}_0$ be the social group with support over $[w', \underline{w}]$ and $\mathscr{L}_1$ the social group with support over $[\underline{w}, \overline{w}_1]$. Furthermore, let ϕ_0 and ϕ_1 be the group qualities. Let $\underline{\phi}'$ be the lower bound of ϕ given W' and $\mathscr{L}'$. It is for the following argument without loss to assume $\phi_0 = \underline{\phi}'$.

Fix a given $\mathscr{L}'$. Now construct an alternative measure $\mathscr{L}^\epsilon$ with $\mathscr{L}'(W') = \mathscr{L}^\epsilon(W')$ such that almost all agents in $[w', \underline{w}]$ are 'close' to $\underline{w}$. In particular, let $\mathscr{L}^\epsilon$ be such that $\mathscr{L}^\epsilon([\underline{w} - \epsilon, \underline{w}]) = \mathscr{L}^\epsilon([w', \underline{w}] - \epsilon$. This is defined for $\epsilon < \min\{\underline{w} - w', \mathscr{L}'([w', \underline{w}])\}$. Note that as $\epsilon \to 0$, the measure approaches the (scaled) dirac measure.The measure might influence $\underline{\phi}'$ but since the interval is fixed, we know the quality is bounded by some $\underline{\phi}$ for all measures.

If $\phi_1 \leq \phi_0$, no set of prices exists that would sustain $\mathscr{I}'$ as an equilibrium. Suppose instead $\phi_0 < \phi_1$. Given that almost all agents are within $[\underline{w} - \epsilon, \underline{w}]$, the rank difference $r_0(\underline{w} - \epsilon) - 0$ is by construction ϵ meaning that $r_0(\underline{w}) - r_0(\underline{w} - \epsilon) = 1 - \epsilon$. For ϵ small enough, $u(\underline{w}, \phi_1, 0) - u(\underline{w} - \epsilon, \phi_1, 0)$ is arbitrarily close to 0. But note that $u(\underline{w}, \phi_0, 1) - u(\underline{w} - \epsilon, \phi_0, 0)$ does not converge to 0 as for every w and $r \in [0,1)$, $u_r(w, \underline{\phi}, r) > 0$. As $\epsilon \to 0$, agents at $\underline{w} - \epsilon$ have a strict incentive to join g_1 instead of g_0. The partition is not incentive compatible. As we used the lower bound $\underline{\phi}$ to show this, the result is independent of how $[w', \underline{w}]$ and is partitioned. □

Proof of Proposition 10:

Proof. I demonstrate that splitting a group into two groups is less beneficial to a monopolist under status concern. Since every refinement can be written as an iteration of such splits, this suffices to prove the proposition.

Take any equilibrium group provision $m(w)$ and social group $\mathscr{L}_g$ with support over $[\underline{w}_g, \overline{w}_g]$. Recall that p denotes the price vector arising from $m(w)$.
It follows from the discussion on pricing that the price a monopolist can charge for agents in g is $p_g = u(\underline{w}_g, \phi_g) + v(\underline{w}_g, 0) - \kappa$ where κ is either the stand-alone utility $\underline{u}$ or, if there is a group 'below', κ is the utility an agent of type $\underline{w}_g$ would obtain if he reported a type that would have him assigned to the group below i.e. $U(\underline{w}_g, m(\underline{w}_g - \epsilon), \mathscr{L}_A)$ for an arbitrarily small $\epsilon > 0$. Since any split of g does not affect the utility and prices in the groups below, we can treat κ as a constant. Any refinement of g affects the price at $\underline{w}_g$ only through changes in quality.

Suppose the monopolist offers a finer provision $m'(w)$ such that g is split into two groups $[\underline{w}_l, \overline{w}_l]$ and $[\underline{w}_h, \overline{w}_h]$ where $\underline{w}_l = \underline{w}_g$ and $\overline{w}_h = \overline{w}_g$ with prices p'_l and p'_h.

It follows from incentive compatibility that all prices higher than p_g in the original p need to be adjusted to make the new provision incentive compatible. The effect on revenue is determined by the change of all prices $p'_i \geq p'_l$ noting that for every price in p with $p_i \geq p_g$, there is a corresponding price p'_i in p' and there is one additional price p'_h that does not map from any previous price. We decompose these into the effect caused by $p'_l - p_g$ which affects all prices $p_i \geq p_g$ and the effect of introducing a new price p'_h, which increases all prices $p_i > p_g$ by $p'_h - p'_l$.

It follows from the discussion on incentive compatibility that $\phi_l < \phi_h$ and $p'_l < p'_h$ otherwise $m'(w)$ could not be an equilibrium provision. Because of the separability, we can conclude that $p'_l - p_g$ is the same for preferences with and without status concern. It is equal to $u(\underline{w}_g, \phi_l) - u(\underline{w}_g, \phi_g)$ noting that, for preferences with status concern, the v component cancels. All prices $p(w)$ in the initial provision greater than p_g, if any, have to adjust by the same difference $u(\underline{w}_g, \phi_l) - u(\underline{w}_g, \phi_g)$ in the new provision.

Furthermore, the 'split' adds an 'additional' price p'_h and all prices in the initial provision with $p(w) > p_g$ need to further adjust by $p'_h - p'_l$ for the new provision. It follows from incentive compatibility that $p'_h = p'_l + u(\underline{w}_h, \phi_h) - u(\underline{w}_h, \phi_l) - \left(v(\underline{w}_h, 1) - v(\underline{w}_h, 0)\right)$. The difference to the price these agents paid under p is equal to:

$$p'_h - p_g = p'_l - p_g + u(\underline{w}_h, \phi_h) - u(\underline{w}_h, \phi_l) - \left(v(\underline{w}_h, 1) - v(\underline{w}_h, 0)\right) \tag{2.17}$$

As shown before, $p'_l - p_g$ is the same with and without status concern while $v(\underline{w}_h, 1) - v(\underline{w}_h, 0)$ is strictly positive under status concern and 0 without. The change in revenue from splitting any group is strictly larger without status concern. □

Proof of Corollary 10.1:

Proof. The argument follows that in the proof of Proposition 10. We first observe that $p'_l - p_g$ does not depend on α as the effect of a change in quality on the utility of a type-$\underline{w}_l$ agent does not depend on α. However, the change in price paid by every agent in $[\underline{w}_h, \overline{w}]$ as defined in Equation (2.17) needs to be restated as:

$$p'_h - p_g = p'_l - p_g + u(\underline{w}_h, \phi_h) - u(\underline{w}_h, \phi_l) - \alpha\left(v(\underline{w}_h, 1) - v(\underline{w}_h, 0)\right)$$

This is decreasing in α noting that $p'_l - p_g$ is unaffected by α. The change in revenue is larger for smaller α as the price difference that makes the split incentive compatible is decreasing in α. □

Proof of Result 10.2:

Proof. Both $\mathscr{I}$ and $\mathscr{I}'$ are equilibrium group structures. Let again p and p' be the price vectors generated from $p(w)$ and $p'(w)$. I denote the lowest price in p as p_1 and the lowest in p' as p_0 since p' has one more price than p.
For the proof, I distinguish between the price vectors under status concern (p, p') and the prices inducing the same structure without status concern (q, q').

The total gain in revenue from offering the full-participation group structure $\mathscr{I}'$ under status concern is:

$$p'_0 - (p_1 - p'_1)[1 - F(w_1)]$$

and equivalently for preferences without status concern.
We observe that

$$p'_1 = p_1 - \left[U(w_1, m'(\underline{w}), \mathscr{L}_g) - U(\underline{w}, m'(\underline{w}), \mathscr{L}_g)\right]$$

i.e. the intra-group utility difference for the lowest group under m'. (Recall that $U(w_1, m'(\underline{w}), \mathscr{L}_g) = u(w_1, \phi_0, 1)$ since $r(w)$ is determined by the true not the reported type). Furthermore, $p'_0 = u(\underline{w}, \phi_0, 0) - \underline{u}$.

The prices without status concern are accordingly:

$$q'_1 = q_1 - \left[U^q(w_1, m^{q\prime}(\underline{w}), \mathscr{L}_A) - U^q(\underline{w}, m^{q\prime}(\underline{w}), \mathscr{L}_A)\right]$$

and $q'_0 = u(\underline{w}, \phi_0, r) - \underline{u}$.

As $u(w, \phi, 1) - u(w, \phi, 0) > 0$, it follows from complementarity that the intra-group utility difference is larger under status concern meaning:

$$\left[U^q(w_1, m^{q\prime}(\underline{w}), \mathscr{L}_A) - U^q(\underline{w}, m^{q\prime}(\underline{w}), \mathscr{L}_A)\right] < \left[U(w_1, m'(\underline{w}), \mathscr{L}_g) - U(\underline{w}, m'(\underline{w}), \mathscr{L}_g))\right]$$

and thus

$$q'_0 = u(\underline{w}, \phi_0, r) - \underline{u} \geq u(\underline{w}, \phi_0, 0) - \underline{u} = p'_0$$

we can conclude that

$$p'_0 - (p_1 - p'_1)[1 - F(w_1)] < q'_0 - (q_1 - q'_1)[1 - F(w_1)]$$

The result follows. □

Proof of Result 10.3:

Proof. Suppose a monopolist offers $m(w)$ instead of $m'(w)$ with prices $p(w)$ and $p'(w)$ respectively. Incentive compatibility requires that the price difference at the lowest cut-off is:

$$p_1 - p'_1 = u(w_1, \phi_1, 0) - u(w'_1, \phi'_1, 0)$$

Since ϕ is bounded, we can set $\phi_h = \overline{\phi}$ and $\phi_l = \underline{\phi}$. If $u(\overline{w}, \phi_h, 0) - u(\underline{w}, \phi_l, 0) = \epsilon$, then this implies that $u(w_1, \phi_1, 0) - u(w_1', \phi_1', 0) < \epsilon$. For any $\mathscr{I}$, incentive compatibility requires:

$$p_1 - p_1' < \epsilon$$

For any higher quality group g, $p_g - p_g' < p_1 - p_1' < \epsilon$. This is due to the fact for incentive compatibility, all such prices increase by $p_1 - p_1'$ but, at the same time, need to be reduced by $u(w_2, \phi_1, 1) - u(w_2, \phi_1', 1) > 0$.
The effect on profits is thus bounded above by

$$\epsilon[1 - F(w_1)] - [F(w_1) - F(w_1')]p_1'$$

As $\epsilon \to 0$, this becomes negative. Excluding agents with types in $[w_1', w_1)$ is not optimal for the monopolist. As a final step, we show that welfare is not bounded by this.
The welfare increase can be positive as for all $r(w) > 0$:

$$u(w, \phi, r(w)) - u(w, \phi', r(w)) > u(w, \phi, 0) - u(w, \phi', 0)$$

and it is thus feasible that

$$\int_{w_1}^{w_2} \left(u(w, \phi_1, r_1(w)) - u(w, \phi_1', r_1(w)')\right) dF(w) > \int_{w_1'}^{w_1} \left(u(w, \phi_1', r_1(w)') - \underline{u}\right) dF(w)$$

The increase in w_1' weakly lowers the rank of all agents in $[w_1, w_2]$ but if $\phi_1 - \phi_1'$ is sufficiently large, this is beneficial for almost all agents in that group. The aggregate welfare of all other groups is unchanged. The effect of excluding more agents is thus either negative or positive while for sufficiently small ϵ, excluding more agents always lowers profits. □

Proof of Result 10.4:

Proof. Suppose there is a full participation equilibrium with the lowest price $p_1 > 0$. If the price vector arising from the market provision is p, some firm i can offer a $p_i(w)$ such that the price vector for the corresponding provision is $p - \epsilon$ for some $\epsilon > 0$. Since all the price differences between groups remain the same, if the equilibrium provision $\mu(w)$ induced the group structure $\{\mathscr{L}_g\}$ then this is also an equilibrium under the new prices. Since $\epsilon > 0$, all agents are strictly better off. μ' must be such that the group structure is identical but all agents pay the lower price according to $p - \epsilon$. The firm offering these prices would capture the entire market. For small enough ϵ, this increases revenue for any firm not serving the entire market before. □

Proof of Result 10.5:

Proof. Take any equilibrium, full-participation group structure with social groups $\{\mathscr{L}_g\}_{g\in G}$ and prices p. Suppose there are $K \geq 2$ active groups. Suppose further that p is normalized so that the lowest price $p_1 = 0$ - this is always possible for full participation group structures. The lowest price vector p' that maintains the same group structure and achieves 0 profits is such that:

$$p' = p - \sum_{2}^{K} p_k \mathscr{L}_k(W)$$

If N firms are serving the market, then they must all make 0 profits. If not, then there exists some firm providing active groups G^n such that

$$\Pi^n = \sum_{G^n} p'_k \mathscr{L}_k(W) > 0$$

Another firm can offer prices $\tilde{p} = p - \epsilon \Pi^n$ for some $\epsilon > 0$. This is a Pareto-improvement to all agents since the group structure does not change but prices are lower. For ϵ small enough, it strictly increases profits for this firm. This implies at least the lowest active price $p'_1 < 0$ in equilibrium.

Suppose the number of active social groups in $a \geq 2$. Then the number of non-negative active prices is $a^- \geq 1$. Note that if there is no such price, at least one firm would make positive profits since the number of active groups is $K \geq 2$ and so there is at least one strictly positive price. If now $N > a^-$, then there is at least one firm providing only positively priced groups which, by the previous argument, cannot be an equilibrium.

Finally, suppose a^+ is the number of non-negative active prices. (Note that either $a^+ + a^- = a$ or $a^+ + a^- = a + 1$). If $N > a^+$ i.e. the number of firms is greater than the number of strictly positive active prices, then at least one firm must provide only negatively priced active groups. This firm makes negative profits and has a strict incentive to raise prices. This cannot be an equilibrium. □

Proof of Result 10.6:

Proof. Suppose $\mathscr{I} = \{w_1, ..., \overline{w}\}$ is an incentive compatible group structure under p and $\mathscr{I}' = w_0 \cup \mathscr{I}$ with $w_0 < w_1$ is incentive compatible under p'. We note that under p', all active prices p'_i with $i \geq 1$ are less than p_i. To exclude agents in the interval $[w_0, w_1)$, it must be that $p_1 = U(w_1, \mu(w_1), \mathscr{L}_A) + p_1 - \underline{u}$ - an agents with w_1 is indifferent between joining and not joining. Under p', we need that agents with w_0 are indifferent. Incentive compatibility requires that at the cut-off $U(w_1, \mu(w_1), \mathscr{L}_A) = U(w_1, \mu(w_0), \mathscr{L}_A)$

but as utility is increasing in rank, $U(w_1, \mu(w_0), \mathscr{L}_A) > \underline{u}$. It must be that $p_1' < p_1$. Furthermore, since all groups with higher types are unchanged, $p_i - p_i' = p_1 - p_1' > 0$. Since almost all agents in $[w_0, w_1)$ are strictly better off and all agents in higher quality groups pay a lower price, this is a Pareto improvement. Since it allows one firm to capture the entire market, it follows from previous arguments that there is a strict incentive for at least one firm to offer p' if all other firms offer p. $\mathscr{I}$ is not an equilibrium group structure. □

2.8.2 B: Comparison Table

We first observe that in all cases, the price difference between two adjacent groups is lower under status concern as $u(w, \phi_h, 0) - u(w, \phi_l, 1) < u(w, \phi_h, r) - u(w, \phi_l, r)$.

Case 1: $\phi_l < \phi_g$ and $\phi_h < \phi_g$
We observe that as ϕ and r are complements, $p_g - p_l' = u(\underline{w}_l, \phi_g, 0) - u(\underline{w}_l, \phi_l, 0)$ is less than $u(\underline{w}_l, \phi_g, r) - u(\underline{w}_l, \phi_l, r)$ for the reference utility at $1 > r > 0$, which is the price change under status concern. Since this leads to a drop in revenue, the effect is weaker under status concern. Furthermore, $u(\overline{w}_h, \phi_g, 1) - u(\overline{w}_h, \phi_h, 1)$ is larger than $u(\overline{w}_h, \phi_g, r) - u(\overline{w}_h, \phi_h, r)$ which has a greater positive effect on revenue if $\overline{w}_h \neq \overline{w}$ (i.e. there is at least one additional active group with higher quality).

Case 2: $\phi_l < \phi_g$ and $\phi_h > \phi_g$
The effect for the low quality group is as in Case 1. For the top group, we note that the increase in quality has a negative effect on revenue at $\overline{w}_h$ if there are groups with higher types. This is because $u(\overline{w}_h, \phi_h, 1) - u(\overline{w}_h, \phi_g, 1)$ is larger than $u(\overline{w}_h, \phi_h, r) - u(\overline{w}_h, \phi_g, r)$. This difference reduces revenue as the price difference to groups above has to fall by that amount. If the weighted difference for the higher p_l' under status concern outweighs the lower prices $p_i' > p_l'$, then the increase in revenue is larger under status concern and vice versa. The overall effect is thus ambiguous.

Case 3: $\phi_l > \phi_g$ and $\phi_h > \phi_g$
The effect of the change $\phi_h - \phi_g$ is as above and has a more negative impact under status concern. To reach the conclusion that the increase in revenue is less under status concern, we note that p_l' must increase by less under status concern since $u(\underline{w}_l, \phi, 0) - u(\underline{w}_l, \phi_g, 0)$ is less than $u(\underline{w}_l, \phi, r) - u(\underline{w}_l, \phi_g, r)$ for any $\phi > \phi_g$. We conclude that under status concern $p_l' - p_g$ is lower, $p_h' - p_g$ is lower and the difference $p_i' - p_i$ in all prices $p_i > p_g$ is less. Under status concern, revenue increases less (decreases more) when offering p' instead of p.

Case 4: $\phi_l > \phi_g$ and $\phi_h < \phi_g$

This would imply $\phi_l > \phi_h$ which is ruled out by Corollary 3.1.

Chapter 3

The Benefits of Being Misinformed[1]

3.1 Introduction

We commonly observe situations in which people take actions in line with hypotheses that seem to be contradicted by empirical evidence. Furthermore, these actions often produce additional information that should highlight the initial mistake. For example, a significant number of people refuse even essential vaccinations despite the very strong evidence of their benefit and despite the measurable increase in outbreaks of the related disease as a consequence of this refusal.[2] For this purpose, we revisit the comparison of experiments as in Blackwell (1951) under the assumption that information processing is not always flawless and might be impeded by systematic mistakes. We do this in a setup that captures the fundamentals of Bayesian updating and its consequences on utility: First an agent takes an action that directly affects his payoff but also provides information about the payoff relevant state of the world. The agent then takes another action before payoffs are realized. Motivated by the psychological and experimental literature on beliefs and perceptions, the information processing can be 'imperfect' in two ways: an agent might initially hold an incorrect prior affecting the type of signals he receives and the agent might misinterpret the signals.[3]

We find that both types of biases by themselves are welfare decreasing. The potential welfare loss can be ordered according to the magnitude of the bias. Next, we provide necessary and sufficient conditions under which a given binary ranking of action profiles can be reversed by a perception bias. Building on these findings, we can show that it is not always true that adding another type of bias to a pre-existing one makes

[1] This chapter is joint work with Marcus Roel.
[2] See, for instance, Wallace (2009).
[3] For example: Bruner and Potter (1964), Darley and Gross (1983), Fischoff, Slovic and Lichtenstein (1977) and Lichtenstein, Fischoff and Lawrence (1982).

an agent even worse-off. In particular, if the agent tends to misperceive signals, then pushing the agent's prior further away from the truth can beneficial if it causes the agent to prefer an action whose signal is less ambiguous and thereby less likely misinterpreted. Our setting thus provides a novel channel for the long known observation by Hirshleifer (1971) that information may not always be welfare improving.
On the other hand, when the agent's prior is incorrect and signals are useful in the sense that they inform decision-making, the agent will always be worse off from misperception. Only in the extreme case, where it's optimal for an agent to always take a fixed action irrespective of any signal, increasing the agent's degree of misperception can be beneficial. It follows that no straight-forward welfare ordering can be established if both types of misperception are present. We further explore the implications of an agents awareness over the bias. We provide conditions under which sophistication regarding the misperception of signals can help or harm the agent's welfare.

To provide an illustration, let's look at a student who decides whether to become an entrepreneur. The literature highlights that entrepreneurship is ex-ante not very profitable. Landier and Thesmar (2009) show that entrepreneurs tend to be overly optimistic with respect to their company's future growth prospects. Moreover, their optimism is positively correlated with higher-education, and more optimistic entrepreneurs tend to choose more short-term debt. Suppose the payoffs from being an entrepreneur depend on the student's ability. The student can either quit university immediately in order to start his business, or wait and finish his degree first. While graduating is inherently useful, it also provides him with a signal about his abilities. Finally, when he starts his business, the student also has to decide how much to borrow.
A student who is very overconfident in his abilities will find it optimal to start his business immediately, yet may not be overconfident enough to borrow much. In contrast, a slightly less overconfident agent will pursue the more 'sensible' path of completing his studies. When the student also suffers from biases in perception, he may (a) misinterpret the signal as confirming his superior ability or (b) over-interpret the signal strength. As a result, the student may come out of university feeling even more confident in his ability, which can persuade him to start a business as well as taking on larger risks, i.e. a bigger amount of debt. While both businesses fail with equal probability, the initially less over-confident agent may end up being worse off in expectation. The more over-confident agent managed to avoid the signal and thereby was not in a position to misinterpret it.

The chapter is organized as follows: In Section 2, we summarize the relevant literature.

This is followed by a basic description of our model in Section 3 and the characterization of the unbiased choice problem in Section 4. Section 5 introduces biased perception and biased priors into the model and relates our setting to Blackwell (1951). This is followed by our main results in Section 6. We conclude with a short investment example in Section 7 and a final discussion in Section 8.

3.2 Literature

Blackwell (1951) formalizes when an information experiment is more informative than another. Marschak and Miyasawa (1968) transfer these statistical ideas to the realm of economics. The key conclusion is that no rational decision maker would choose to 'garble' his information, i.e. voluntarily introduce noise into his experiments.

Having more information, however, may not always be better in an economic settings. Hirshleifer (1971) highlights that public information may destroy mutually beneficial insurance possibilities. The recent behavioral literature takes this idea further: Papers on overconfidence have shed light on how holding incorrect beliefs can be useful. These benefits arise from strategic interaction between agents. Ludwig, Wichardt and Wickhorst (2011) show that overconfidence can improve the agents relative and absolute performance in contests by inducing higher efforts. De La Rosa (2011) studies the effect of overconfidence on incentive contract in a moral hazard problem. When agents overestimate the probability of success as well as their marginal contribution to success, it becomes easier for the principal to induce effort. It turns out that the efficiency gains from slight levels of overconfidence can improve the agent's welfare.

Carrillo and Mariotti (2000) show that Blackwell garbling of information may increase the current self's payoff when individuals are time-inconsistent. Benabou and Tirole (2002) develop this idea further. In their model, a time-inconsistent agent can take on a project with deferred and uncertain benefits but immediate losses. Due to time-inconsistency a time-0 self prefers that the project is undertaken in the next period, but anticipates a lack of motivation by her next period's incarnation. Having the opportunity to either perfectly learn the true success rate or to stay uninformed, she may prefer to stay uninformed if her prior motivates the time-1 agent to work. While these two papers are phrased as a decision problem, an agent with time-inconsistency plays a game among his different selves, which is the fundamental driver of these results.

Other papers analyze how various behavioral shortcomings can be improved upon by overconfidence. These papers aim to provide a motivation why overconfidence exists in the first place. In Compte and Postlewaite (2004), failing to recall past failures can improve welfare as it counteracts the fear of failure. Brunnermeier and Parker (2005) highlight that agents prefer to hold incorrect, too optimistic beliefs about the future, when they derive immediate benefits from these expectations.

Our paper is also related to the enormous literature on mistakes in updating expectations. Given our focus on how basic biases in information processing interact, we shall keep references to specific biases short. In our model, agents have the tendency to misperceive information, which represents a generalisation of the confirmatory bias in Rabin and Schrag (1999). In their paper, they show how confirmatory bias, the tendency to misinterpret new information as supportive evidence for one's currently held hypothesis, can not only lead to overconfidence in the incorrect hypothesis, but even cause someone to become fully convinced of it.

There have been many studies that suggest people hold incorrect beliefs. On average, people tend to have unrealistically positive views of their traits and prospects. To mention a few, see Weinstein (1980) for health and salaries, Guthrie, Rachlinski and Wistrich (2001) for rates of overturned decisions on appeal by judges, and Fischoff, Slovic and Lichtenstein (1977) as well as Lichtenstein, Fischoff and Lawrence (1982) for estimates of ones' own likelihood to answer correctly. Recent papers include Landier and Thesmar (2009), for entrepreneurs, as well as Malmendier and Tate (2005) linking CEO overconfidence to a higher likelihood of pursuing risky actions, for instance acquisitions. Benoit and Dubra (2011) argue that a lot of the empirical evidence is also consistent with Bayesian updating under correct priors and thus may not demonstrate overconfidence. In response, laboratory experiments robust to this criticism were carried out by Burks et al. (2013), Charness, Rustichini and van de Ven (2014), and Benoit and Moore (2015), again documenting overconfidence.

3.3 The Setting

We consider a two period model with two states of the world $\Omega = \{A, B\}$. The agent has a finite set of actions $X = X_1 \times X_2$. In the first period, he chooses an action from X_1 and subsequently receives a signal about the state. The quality of the signal depends on the action he chooses. He then decides on a second action from X_2 and receive a payoff determined by the action profile as well as the state of the world. The probability that a particular state materializes is $\Pr(\omega = A) = p \in (0,1)$ and $\Pr(\omega = B) = 1 - p$ respectively. Denote the agent's prior belief that the state is A by μ. This belief may or may not coincide with the correct p. Let $u(x_1, x_2|\omega)$ denote the payoff if x_1 is taken at $t = 1$, x_2 at $t = 2$ and the state is ω, with $u(x_1, x_2|\omega) \in \mathbb{R}$.

We make the following assumptions on payoffs which guarantee that in each state there is a unique best action profile:

Assumption A1: There exist action profiles:
(a) $(x_1^A, x_2^A) \in X$ such that $u(x_1^A, x_2^A|A) \geq u(x_1, x_2|A)$ for any $(x_1, x_2) \in X$.

(b) $(x_1^B, x_2^B) \in X$ such that $u(x_1^B, x_2^B|B) \geq u(x_1, x_2|B)$ for any $(x_1, x_2) \in X$.

To avoid trivial scenarios, we exclude cases in which those two profiles have any common component. It will also be implicitly assumed that X only contains actions that are not completely payoff equivalent and ties are broken deterministically. Additionally, payoffs are assumed to be bounded:

Assumption A2: There exists a $K \in \mathbb{R}$ such that $K > \max\{u(x_1^A, x_2^A|A), u(x_1^B, x_2^B|B)\}$.

After period one, an agent receives a private signal $s \in S$. For each action $x_1 \in X_1$, there is a binary set of potential signals $S(x_1) = \{a(x_1), b(x_1)\} \subset S$. The probability distribution over those signals is conditional on the action as well as the state of the world meaning that x_1 defines a probability measure on $S \times \Omega$. Let $\pi(x_1, \omega) \equiv Pr(s = a(x_1)|\omega, x_1)$ and consequently $1 - \pi(x_1, \omega) \equiv Pr(s = b(x_1)|\omega, x_1)$. Let the signal structure be symmetric between states, i.e. $\pi(x_1, A) = 1 - \pi(x_1, B)$, and let signals be weakly informative in the sense that a-signals are more likely than b-signals in state A. We call $(x_1, \{x_a, x_b\})$ an *action profile* where $x_a, x_b \in X_2$ represent the actions taken after an $a(x_1)$ or $b(x_1)$ signal. Such an action profile is said to be *signal-sensitive* if $x_a \neq x_b$. A *simple action profile* is such that $x_a = x_b$ meaning the second-period choice is not conditional on the signal. It is for brevity denoted by (x_1, x_2).

We can think of the first action as an experiment that delivers information about the realized state. The agent can react accordingly and adjust his action in the second period. Notice, however, that an experiment is only useful if it is not too costly. For instance, it could perfectly reveal the state but reduce the attainable utility to an extent that it would be better to choose a noisier experiment.

3.4 Unbiased Choice Problem

We reverse-engineer the agent's decision problem step by step which serves to clarify the later results. First we look at optimal actions in the second period and then focus on the optimal choice of experiments in the first period. The discussion is rather technical and the descriptions necessarily somewhat dry. But it allows for more constructive proofs later. Readers impatient for key results may skip ahead and return to this section as needed for later understanding.

3.4.1 Period 2 Cutoff-Strategy

The expected utility of an action $x_2 \in X_2$ from the second period's perspective depends on the posterior after receiving the signal. Furthermore, it can also depend directly on the action in the previous period. Let $\mu(s)$ be the posterior after receiving a type-s signal. An agent chooses action x_2 at $t = 2$ given a posterior $\mu(s)$ and an action x_1 at $t = 1$ if

$$E\left[u(x_1, x_2)|\mu(s)\right] \geq E\left[u(x_1, x)|\mu(s)\right] \qquad \forall x \in X_2$$

The expected payoff from any given action in period 2 is monotonic in beliefs. Starting from some $0 < \mu < 1$, an increase in μ strictly increases the expected payoff from (x_1^A, x_2^A) and decreases the one from (x_1^B, x_2^B). In fact, for any $(x_1, x_2) \in X$ with $u(x_1, x_2|A) \neq u(x_j, x_k|B)$, the expected payoff is either strictly in- or decreasing in μ. We can order actions at $t = 2$ according to expected payoffs based on $\mu(s)$ - the belief after receiving the signal s. This gives rise to a cutoff-type decision rule.

To illustrate this, consider the period 2 expected payoff of x_2^A given the action x_1^B in period 1 and some posterior $\mu(s)$:

$$\mu(s)\, u(x_1^B, x_2^A|A) + (1 - \mu(s))\, u(x_1^B, x_2^A|B)$$

Similarly, the expected payoff from x_2^B is

$$\mu(s)\, u(x_1^B, x_2^B|A) + (1 - \mu(s))\, u(x_1^B, x_2^B|B)$$

If $u(x_1^B, x_2^A|A) > u(x_1^B, x_2^B|A)$, there exist some $\mu^*(s) \in (0,1)$ such that for all $\mu(s) > \mu^*(s)$, the expected payoff from choosing x_2^A is larger than from x_2^B. Equally, for all $\mu(s) < \mu^*(s)$, the expected payoff from x_2^B exceeds the expected payoff from x_2^A as $u(x_1^B, x_2^A|B) < u(x_1^B, x_2^B|B)$. An equivalent argument shows that the same is true for any other $x_2 \in X_2$ such that $u(x_1^B, x_2|A) > u(x_1^B, x_2^B|A)$. For μ large enough, the expected payoff from x_2 exceeds the one from x_2^B and vice versa for μ small enough. Iterating this argument over all available actions at $t = 2$ and each $t = 1$ action allows us to conclude that the region in which a given x_2 is chosen must be connected. Result 12 in Appendix 3.9.3 shows this formally.

We can now derive the conditions under which a given action is chosen in period 2 for at least some beliefs. For any $x_1 \in X_1$ there is a $x_a \in X_2$ such that $u(x_1, x_a|A) \geq u(x_1, x_k|A)$ for all $x_k \in X_2$. x_a is optimal for $\mu(s) = 1$ and will be chosen if $\mu(s)$ is high enough. The equivalent is true for some action in state B. If the actions are identical across states then this is the best possible action for all $\mu(s)$. Otherwise, there might be different optimal actions for intermediate beliefs.
Consider the threshold belief $\mu_{a,j}$ for which the expected payoff from (x_1, x_a) is exactly equal the one from (x_1, x_j):

$$\mu_{a,j}\, u(x_1, x_a|A) + (1-\mu_{a,j})\, u(x_1, x_a|B) = \mu_{a,j}\, u(x_1, x_j|A) + (1-\mu_{a,j})\, u(x_1, x_j|B)$$

We can rearrange the equation to

$$\frac{\mu_{a,j}}{(1-\mu_{a,j})} = \frac{u(x_1, x_j|B) - u(x_1, x_a|B)}{u(x_1, x_a|A) - u(x_1, x_j|A)}$$

Since by definition $u(x_1, x_a|A)$ is the highest utility in state A for x_1, the equation highlights that $u(x_1, x_j|B) > u(x_1, x_a|B)$ is a necessary and sufficient condition for this threshold to exist. We can then simply order actions according to payoffs in both states and ignore actions that are dominated in both states. Starting from $\mu_a = 1$, for which x_a is the best action, we can compare the potential cutoffs for all actions $x_i \in X_2$ that are not strictly dominated. The action with the highest $\mu_{a,i}$ will be the one chosen for some range of beliefs. We then continue to iterate this process from this action until there is no more action that has a higher payoff in state B.

Result 1: *For every $x_1 \in X_1$, there exists a partition $\mathscr{P}_{x_1}$ of $[0,1]$ such that for every two consecutive elements p_i and p_{i+1} of the partition, there is an action $x_2 \in X_2$ such that $E[u(x_1, x_2)|\mu] \geq E[u(x_1, x)|\mu]$ for all $x \in X_2$ and $p_i < \mu < p_{i+1}$.*

Result 1 follows directly from the previous argument and Result 12. It highlights how the choice in period 2 depends on the posterior, which in turn is determined by the signal. Differences in the posterior are only welfare relevant if they fall in different elements of the partition. For binary signals, there are at most 2 different choices in period 2 and thus potentially 4 different utility outcomes. Since those are pinned down for every $x \in X_1$ by $\mathscr{P}_x$, we can collapse the problem to a comparison of experiments in period 1, fixing the corresponding optimal period 2 choices.

3.4.2 Choice in Period 1

In the first period, the decision maker simply maximizes his expected utility:

$$\max_{x_1 \in X_1} E\left[u\big(x_1, \{x^*_{a(x_1)}, x^*_{b(x_1)}\}\big)|\mu\right]$$

where $\{x^*_{a(x_1)}, x^*_{b(x_1)}\}$ are the optimal period 2 actions following $x_1 \in X_1$. The period 1 choice balances the information value of an experiment as well as the immediate utility derived from it. A very informative action leads to very different posteriors and, keeping in mind the partition, to different actions in period 2. When the experiment is uninformative, the posterior $\mu(s)$ equals the prior μ and falls into the same bracket of the partition.

Taking x_1 as given, let $\mu(s)$ be the posterior after receiving a type-s signal and x^*_s be the corresponding optimal action. We can write the expected utility as

$$\begin{aligned}&[\mu(a)u(x_1, x^*_a|A) + (1-\mu(a))u(x_1, x^*_a|B)]Pr(s=a|x_1,\mu)\\ +&[\mu(b)u(x_1, x^*_b|A) + (1-\mu(b))u(x_1, x^*_b|B)]Pr(s=b|x_1,\mu)\end{aligned}$$

Recall that the probability of receiving an $a(x_1)$ signal in state A is $\pi(x_1, A) \geq \frac{1}{2}$ which is also equal the probability of receiving the signal $b(x_1)$ in state B. Using Bayes' rule, we can rewrite the expected utility expression as:

$$\begin{aligned}\mu \quad & [\pi(x_1,A)u(x_1,x^*_a|A) + (1-\pi(x_1,A))u(x_1,x^*_b|A)]\\ +(1-\mu) \quad & [\pi(x_1,B)u(x_1,x^*_a|B) + (1-\pi(x_1,B))u(x_1,x^*_b|B)]\end{aligned} \tag{3.1}$$

This is the objective function for the utility maximization problem at μ. It is a weighted average of receiving the "correct" signal and thus choosing the correct action and the probability of receiving the "incorrect" signal and thus choosing the action that yields the lower utility in the realized state. A higher informativeness in the sense of $\frac{\pi(x_1,A)}{1-\pi(x_1,A)}$ reduces the likelihood of such a mistake. An agent might then be confident enough in the signals that he chooses actions that have a higher variation between states.

We finish this section with a key property of the unbiased agent's problem. Some of the later results arise from a violation of this.

Result 2: *The maximum expected utility at $t = 1$ as a function of μ is convex in μ.*

The previous discussion is best illustrated with a simple example:

Example 1: In each period, there are three actions $\{x_A, x_I, x_B\} = X_1 = X_2$. The actions x_A, x_B yield uninformative signals with $\pi(x_A, A) = \pi(x_B, A) = 0.5$. Action x_I provides an informative signal with $\pi(x_I, A) = 0.75$. Utilities are symmetric in states such that $u(x_A, x_A|A) = 5 = u(x_B, x_B|B)$, $u(x_I, x_A|A) = u(x_I, x_B|B) = 4$ while all other combinations yield 0 utility. x_I represents a pure information experiment that is only useful because it indicates the true state and thus the appropriate action at $t = 2$.
In this setting, the agent never chooses a combination of x_A and x_B. If the agent takes the information experiment in period 1 then he will choose between x_A and x_B in period 2 depending on whether his posterior is greater or smaller than $\frac{1}{2}$.

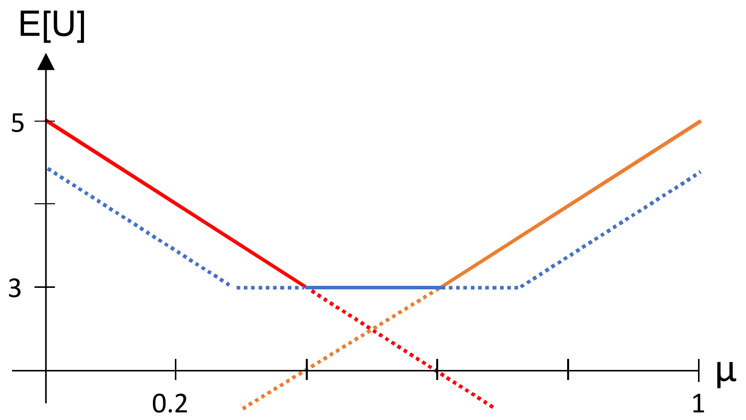

Figure 3.1: Utility frontier

Figure 3.1 illustrates the expected utility outcomes. Actions x_A (orange) and x_B (red) are optimal for extreme enough beliefs. For intermediate beliefs, the informativeness of x_I (blue) is more valuable. The posterior can fall into both partition elements such that x_A or x_B is chosen at $t = 2$ conditional on the signal. His maximum expected utility, indicated by the bold line, is clearly convex in μ.

3.5 Biased Perception

We now turn to our main area of investigation: scenarios in which the agent might misinterpret parts of his environment. In particular, he might not always perceive information accurately or might commit systematic mistakes when observing signals. We call this a *perception bias*. Furthermore, the agent could have a prior that is deviating from the true probability distribution over states. We call this somewhat loosely a *bias in prior*.

The perception bias adds noise to the signal; a form of Blackwell garbling. It weakens the correlation between the signal and the state and thus reduces the information value of experiments in period 1. The agent arrives at incorrect posteriors following the initial signal and may choose suboptimal actions in the second period. The bias in prior might lead to a suboptimal choice of experiment in period 1 with potential consequences for the subsequent choices.

The perception bias could be due to an agent's unwillingness to change their preferred hypothesis or a systematic mistake in interpreting information. Initially, we take no stand regarding the source and exact form of the bias. Unless stated otherwise, we make the crucial assumption that the agent is naive in the sense that he is unaware of this perception problem. He computes his posterior and chooses his actions as if he was receiving signals according to the true underlying structure. In a later section, we discuss how this is related to a setting where agents have no perception bias but simply a distorted view about the accuracy of signals. We argue that they overlap but are not identical. We also explore the implication of a sophisticated agent, who is aware of his perception problem.[4]

To illustrate the setting, imagine the following thought experiment. A doctor orders a medical test. The outcome can be either negative or positive. On top of any inaccuracies of the test itself, suppose there is a certain chance the lab technician enters the incorrect result in the patient's file. The attending doctor never misperceives information if the technician does his job perfectly. However, there will be an information loss if, for instance, the technician mistakenly enters a positive result if the test actually came back negative. It is as if the doctor misperceives the signal. The distortion does not have to be balanced but is independent of the state of the world: the lab technician does not have any knowledge about the truth other than through the test result he observes. If the doctor is unaware of his technician's potential mistake, we refer to him as naive; if he takes his mistakes into account, we call him sophisticated. As mentioned before, we mostly maintain the assumption of a naive doctor. An incorrect prior regarding the condition of the patient might simply lead to an unnecessary test (or failure to conduct a necessary one).

[4] Note that a distinction between a sophisticated and naive agent would be meaningless for a bias in prior.

Let s^t be the signal observed by the technician conducting the experiment x_1, and s the signal received by the doctor. Then

Assumption A3: $\Pr(s, s^t|\omega, x_1) = \Pr(s^t|\omega, x_1) \cdot \Pr(s|s^t, x_1) \qquad \forall \omega \in \{A, B\}$

Define the signal probability distortion function d as a function that converts the true signal probabilities into (potentially different) probabilities that capture the likelihood with which the agent actually perceives the signals. In particular, $d\left((\pi, 1-\pi), S(x_1), \mu\right)$ gives the vector of probabilities for a and b signals for a belief μ, signal structure $S(x_1)$, and the undistorted probability structure $(\pi, 1-\pi)$. We can break this down into two functions, d_a and d_b, one for each type of signal. Let the probability of correctly transmitting an a signal when performing the experiment x_1 be $k_a(x_1, \mu)$, and let $k_b(x_1, \mu)$ be the corresponding probability for a b signal. Notice that the likelihood of mistakes may depend on the prior of the agent as well as the type of signal. We can then write the signal probability distortion function as:

$$d\left((\pi, 1-\pi), S(x_1), \mu\right) = \begin{pmatrix} d_a(\pi, x_1, \mu) \\ d_b(1-\pi, x_1, \mu) \end{pmatrix} = \begin{pmatrix} k_a(x_1, \mu)\pi + (1 - k_b(x_1, \mu))(1-\pi) \\ k_b(x_1, \mu)(1-\pi) + (1 - k_a(x_1, \mu))\pi \end{pmatrix} \tag{3.2}$$

for any $\pi \in [0, 1]$ with the obvious restriction that $d_a(\pi, x_1, \mu) + d_b(1-\pi, x_1, \mu) = 1$. An increase in the *magnitude of the bias* means a reduction of $k_a(x_1, \mu)$, or $k_b(x_1, \mu)$, or both.

Result 3: *The signal probability distortion function d takes the form of equation* (3.2). *It is state independent in the sense that $k_a(x_1, \mu)$ and $k_b(x_1, \mu)$ are state independent.* [5]

For convenience, we can represent the actual experiment as well the distorted one by a Markov matrix. In this notation, an experiment is characterized by a 2×2 Markov matrix K, where element K_{ij} refers to the probability of receiving a signal $s = j$ in state ω_i.

Denote the unbiased, information experiment arising from x_1 by P_{x_1} with the structure:

$$P_{x_1} = \begin{bmatrix} \pi & 1-\pi \\ 1-\pi & \pi \end{bmatrix} \tag{3.3}$$

[5]Notice that the true probability of receiving a certain signal is obviously still conditional on the state. If we want to think of this distortion as confirmation bias along the lines of Rabin and Schrag (1999), with an agent being biased towards hypothesis B and thus misinterpreting a signals as b signals with probability δ, but never the other way around, we set: $k_a(x_1, \mu) = 1 - \delta$, $k_b(x_1, \mu) = 1$.

The agent's perception bias is represented by a Markov matrix $M_{P_{x_1}}$:

$$M_{P_{x_1}} = \begin{bmatrix} k_a(x_1,\mu) & 1-k_a(x_1,\mu) \\ 1-k_b(x_1 i,\mu) & k_b(x_1,\mu) \end{bmatrix} \tag{3.4}$$

The resulting probability structure of the experiment $P_{x_1} M_{P_{x_1}}$ captures exactly the idea behind the signal probability distortion function:

$$P_{x_1} M_{P_{x_1}} = \begin{bmatrix} d_a(\pi,x_1,\mu) & d_b(1-\pi,x_1,\mu) \\ d_a(1-\pi,x_1,\mu) & d_b(\pi,x_1,\mu) \end{bmatrix} \tag{3.5}$$

3.5.1 Blackwell (1951)

We now take a small detour to Blackwell's seminal research on the comparisons of experiments (Blackwell (1951), Blackwell and Girshick (1954)).[6] This digression is useful as it highlights the general principles behind information experiments. We will also see how Blackwell's general results translate to settings where agents are biased - both naive, or sophisticated with respect to their perception mistakes.[7]

In Blackwell's comparison of experiments, an agent compares statistical information experiments. Our setting is slightly more general in the sense is that experiments not only yield information but may also have direct payoff consequences. Upon receiving a signal from an information experiment, the agent maximizes his utility. Denote the value function of experiment P by $V(P)$.

Definition B.1: Let P and Q be two experiments. We say that experiment P is more informative than Q, denoted by $P \supset Q$, if there is a 2×2 Markov matrix M with $PM = Q$.

This definition highlights how the information experiment P is garbled into Q by the matrix M; it is as if noise was added to P to create a less informative experiment Q. From this statistical notion, Blackwell derives his famous result:

Result B1: *Let P and Q be two experiments.* $V(P) \geq V(Q)$ *if and only if there is a* 2×2 *Markov matrix M such that* $PM = Q$.

Simply put, a more informative experiment implies a higher utility. Surprisingly, the

[6]Both are phrased very much in the language of a statistician. For an economist's take on this see Marschak and Miyasawa (1968) who translated the setup into a utility-framework.

[7]For further details about the general setting of Blackwell, please see the Appendix B. All statements generalise to n-signals and n-states. These can also be found in the Appendix.

converse is also true: If one experiment yields a higher expected utility than another, it must be more informative.

In Blackwell's setup the agent is rational. He fully understands the signal generating structure of the two experiments and correctly perceives each signal. The garbling matrix doesn't represent the agent's misperception but is a tool to order the experiments by informativeness.

In order to talk about a Blackwell setting when agents misperceive information, we need a new function representing the agents' expected utility from the experiments. Suppose the unbiased, true information experiment is P but the agent misperceives some signals according to matrix M. As a result, the agent faces an information experiment PM. We denote the expected utility of a naive agent, who is unaware of his perception bias, by $V_n(P, M)$. The first argument always indicates the original signal generating process and the second argument the agent's misperception. This agent believes the information experiment is P and thus chooses his action as if he was rational facing P.

If the agent is sophisticated, that is he is cognisant of his misperception M, we write his utility as: $V_s(P, M)$.

At this point, all tools and notation are developed, the stage is set for agents who are unbiased, sophisticated or naive with respect to their signal processing.

3.6 Biases and Their Implications

We start with the statement that the agent is always worse off with a perception bias than without. This observation follows from the original Blackwell result on the ordering of information experiments.

Result B2: *Take two experiments O and P, with $O \supset P$, and let M be a Markov matrix with $OM = P$. Then $V(O) \geq V(OM) = V_s(O, M) \geq V_n(O, M)$.*

The first inequality shows the welfare effect for sophisticated agents, which follows immediately from the Blackwell result. In the Blackwell setting, the agent has a choice between the original and the garbled experiment. Since it is impossible to reverse one's bias, a biased agent does not actually have this choice and so will always be worse off. The second inequality says that a naive agent is (weakly) worse off than a sophisticated given the same bias. This arises from two observations. One, both the sophisticated and the naive receive the same signals with the same likelihoods. Two, for each signal, the sophisticated is aware of its lower information value, which allows him to make better decisions. It follows that:

Proposition 1: *Take any information experiment P. A sophisticated agent's utility from P is (weakly) larger than the naive agent's utility, but weakly worse than an agent that doesn't suffer from any misperception.*

For our specific choice problem, a stronger statement is possible. We say that the magnitude of the bias increases if k_a, k_b, or both, decrease, i.e. when the off-diagonal of the garbling matrix M increases.

Result 4: *Any signal distortion function where $d_a(\pi(x_1,\omega), x_1, \mu) \neq \pi(x_1,\omega)$ for some $\omega \in \Omega$ (weakly) reduces the expected utility if $x_1 \in X_1$ is chosen at $\mu = p$. It strictly reduces the expected utility if the action profile is signal sensitive.*
The welfare loss is increasing in the magnitude of the bias.

Any perception bias is bad news for an agent who conditions his choices on signals. It leads to more "mistakes" in period 2 choices. As the misperception increases, i.e. d_a moves further away from π, the mistakes increase and the expected utility falls. The welfare consequences of a perception bias can thus be ordered for a given direction of bias.

Looking at this from the perspective of the distortion function d, the following Result 5 shows that we can interpret an increase in the magnitude of the bias as a reduction in the differences of the signal probabilities across states. As the bias increases, the likelihood of receiving a given signal converges across states. Given the previous result, this leads to a welfare loss.

Result 5: *An increase in the magnitude of the bias for some x_1 and μ leads to a decrease in*

$$d_a(\pi(x_1,\omega), x_1, \mu) - d_a(1-\pi(x_1,\omega), x_1, \mu)$$

Using the earlier example, we can illustrate the welfare consequences of a bias for a signal sensitive action profile:
Example 2: Take the setting from Example 1 and suppose that the agent has a confirmation bias towards B. This implies that he perceives $b(x_I)$ too often, i.e. $k_a(S(x_I),\mu) = 1-\delta$, but never misperceives $a(x_I)$, i.e. $k_b(S(x_I),\mu) = 1$. The following graph plots the utility frontier for $\delta = 0.2$ which increases the likelihood of receiving a b signal in state B from $1-\pi(x_I,B)$ to $1-\pi(x_I,B)+0.2\pi(x_I,B)$ and in state A accordingly.
The utility frontier, as seen in Figure 3.2, is strictly lower for any interval in which x_I is actually chosen and identical otherwise. It does not affect the expected utility for more extreme priors, even if x_I was chosen as the posterior will always remain in the same partition regardless of the signal realization. In contrast, the effect is strictly negative for the interval in which the agent conditions the choice on the signal.

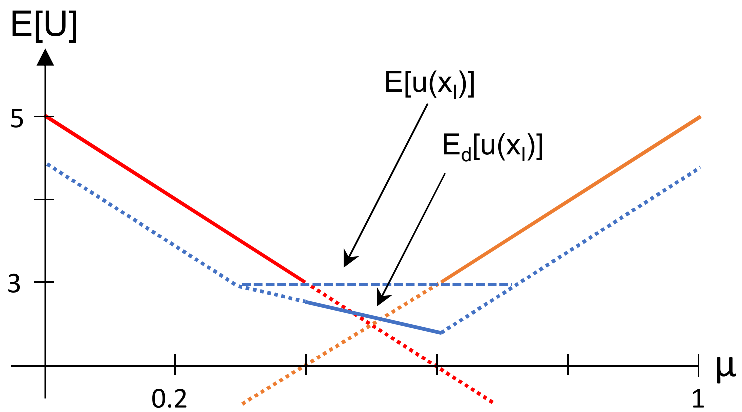

Figure 3.2: Utility frontier with confirmation bias

3.6.1 Biased Prior

We now introduce an alternative source of suboptimal choice: the agent holds an incorrect prior, namely $\mu \neq p$. The prior may be wrong as a consequence of a variety of behavioral and or updating mistakes in the past. If $\mu = p$ the agent will take optimal choices. If, however, $\mu \neq p$, then the true expected utility does not necessarily match the one evaluated by μ.

Equation (3.1) illustrates how a different μ can lead to different choices. An inaccurate μ puts too much weight on one of the states and thus favours actions appropriate for that state. As it realizes with a different probability p, the choice might be suboptimal. It follows:

Result 6: *For any $\mu \geq p$, expected utility is (weakly) decreasing in μ. Equivalently, for $\mu \leq p$ expected utility is increasing in μ.*

3.6.2 Interaction of Biases

So far, we have learned that, in isolation, both biases have a negative effect on outcomes. Additionally, an increase in the magnitude of the biases exacerbates the problem. One might expect that the interaction of both reduces expected utility even further, a 'double whammy'. However, it turns out that this is not the case. To the contrary, an agent with both, a biased perception and a biased prior, can be better off than an agent who only suffers from one of the two. The fundamental reason behind this is that the perception bias shuffles the otherwise straight-forward ranking of experiments at different priors to a different degree. We analyse both directions, i.e. fixing one bias and increasing the other - and determine when a welfare ordering can be established.

Adding a Biased Priors to a Biased Perception

In this section, we assume that the agent's perception is biased 'from the start'. This can either be specific to the experiment - i.e. some experiments generate signals that are more likely to be misperceived than others - or the same across all experiments. We investigate how a bias in prior interacts with the misperception. Breaking the established pattern, we first present an example and then generalize this insight.

Example 3: Reconsider the setting of Example 2 with an agent exhibiting confirmation bias. If the agent is unaware of the bias, the utility frontier is not convex in μ. This arises due to the merely perceived indifference between action x_B and x_I - given optimal period 2 actions - at $\mu = 0.4$. Both achieve an expected utility of 3. The "true" expected utility from x_I is only 2.88 because the agent perceives $b(x_I)$ too frequently and thus takes the period 2 action x_B more often in state A than he would without the bias. The agent should strictly prefer action profile (x_B, x_B). The equivalent situation applies to (x_A, x_A) and μ close to 0.6.

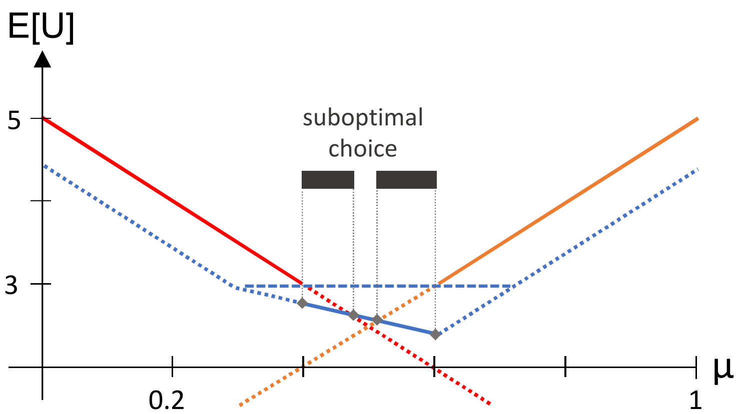

Figure 3.3: Non-convex utility frontier

Now consider the following: if the true p is just below 0.6, then the actually attainable expected utility of an agent with any $\mu > 0.6$ is strictly higher compared to his expected expected utility with $\mu = p$. The choice (x_A, x_A) which would be suboptimal for a fully rational agent, actually yields a strictly better outcome than the "rational choice" x_I. The agent with a biased prior and perception bias fares better than an agent with only perception bias, because his suboptimal choice has a signal structure that is less (or not at all) affected by the perception bias.

To understand the conditions under which such a situation can arise, we introduce the following relation between action profiles:

Definition D1 [Worst-case dominance]: An action profile $(y_1, \{y_a, y_b\})$ *worst-case dominates* $(x_1, \{x_a, x_b\})$ at $\mu \in (0,1)$ if

$$\min_{s\in\{a,b\}} \mu u(x_1, x_s)|A) + (1-\mu)u(x_1, x_s|B) < E[u(\mathbf{y})|\mu].$$

In words, worst-case dominance compares an action profile **y** to the scenario where the agent gets the 'worst' signal realisation for the profile **x**; meaning the signal that leads him to make the worst choice in his contingent plan $\{x_a, x_b\}$ evaluated by the expected utility. Clearly, this worst signal depends on the prior. If the prior is sufficiently close to 1, the worst signal is b and for a prior close to 0, the worst signal is a. Notice that **y** may or may not be preferred to **x** when worst-case dominance holds.
The usefulness of this definition lies in the fact that the concept is independent of the state of the world and enables us to identify which action profiles can be made worse relative to another profile.

Result 7: *For any signal sensitive action profile* **y** *that worst-case dominates* **x** *at* $\mu \in [0,1]$*, there exists a distortion* d *with associated distortion matrices* M_x*,* M_y *such that*

$$E_d[u(\mathbf{x})|\mu] < E_d[u(\mathbf{y})|\mu].$$

If an action profile is worst-case dominated, it is possible to create a distortion that would rank another profile strictly higher in terms of expected utility. Nevertheless, the result is quite weak since it might require a distortion that targets a specific type of signal: the signals $S(x_1)$ generated by experiment x_1. Depending on the scenario, this is hard to justify in practice in a literal way. But we can think of it as a particular experiment yielding harder-to-interpret results while another experiment has very straightforward outcomes. In this context, worst-case dominance tells us something about how the ranking of alternatives can change when one experiment generates very ambiguous signals. We can now state the result that tells us the precise condition when adding a bias in prior can make an already perception-biased agent better off.

Proposition 2: *For any signal sensitive action profile* **x** *that is chosen at* μ *and any action profile* $\mathbf{y} \neq \mathbf{x}$ *that is chosen at* $\mu' \neq \mu$*, there exists a distortion* d *with associated distortion matrices* M_x, M_y *such that* $E_d[u(\mathbf{y})] > E_d[u(\mathbf{x})]$ *when* $p = \mu \neq \mu'$ *if and only if* **y** *worst-case dominates* **x** *at* μ*.*

At least for a relatively unrestricted class of distortions d, we can always find one that makes an agent with an incorrect belief $\mu' \neq p$ better-off than an agent with a correct belief as long as there is sufficient potential variation in payoffs. The next corollary shows that this does not hinge on distorting only one experiment. With a notion very similar to worst-case dominance, we can extended the result to a setting where all signals are distorted equivalently. In the context of our example of a lab technician that might enter incorrect results, the tendency to make such a mistake is now equalized across tests rather than one test being more prone to transcription errors than another.

Definition D2 [Simple dominance]: An action profile $(y_1, \{y_a, y_b\})$ *simply dominates* $(x_1, \{x_a, x_b\})$ at $\mu \in (0,1)$ if

$$\mu u(x_1, x_a)|A) + (1-\mu) u(x_1, x_a|B) < \mu u(y_1, y_a)|A) + (1-\mu) u(y_1, y_a|B)$$

or

$$\mu u(x_1, x_b)|A) + (1-\mu) u(x_1, x_b|B) < \mu u(y_1, y_b)|A) + (1-\mu) u(y_1, y_b|B)$$

or both.

Definition D2 breaks each action profile down into the simple profiles that can be derived from it and compares their expected utility. While simple dominance is not directly comparable to worst-case dominance, the idea is similar. Simple dominance compares the expected utility at μ of each of the two simple profiles derived from the action profiles that are being contrasted. Worst-case dominance compares the expected utility of an action profile with the expected utility of the worst simple profile generated from the action profile it is contrasted against.
Returning to our example, simple dominance compares two treatment methods fixing a particular test result without taking into account the information value. If one method is inferior in at least one of those cases, we can find a type of mistake the lab technician could commit equally across those tests such that the treatment optimal in the absence of mistakes is inferior if the mistake is present. In contrast, worst-case dominance compares a test under ideal conditions with one where the lab technician commits the worst possible type of mistake from an ex-ante perspective.

Corollary 8: *For any signal sensitive action profile* $\mathbf{x}$ *that is chosen at* μ *and any action profile* $\mathbf{y} \neq \mathbf{x}$ *that is chosen at* $\mu' \neq \mu$*, there exists a distortion* d *with associated matrices* $M_x = M_y$ *such that* $E_d[u(\mathbf{y})] > E_d[u(\mathbf{x})]$ *when* $p = \mu \neq \mu'$ *if and only if* $\mathbf{y}$ *simply dominates* $\mathbf{x}$ *at* μ*.*

Proposition 2 and Corollary 8 provide a tight characterisation for utility reversals, accounting for first period action's direct utility effect. The reason behind these results is that the naive agent has an incorrect view about the information value from x_1 relatively to some alternative y_1. This is captured in the following result:

Result B3: *: Take two experiments O and P, with perception biases M_O, M_P. For the naive agent, neither the ordering (according to informativeness) of O and P nor the ordering of OM_O and PM_P determines, or is determined by the agents expected utility ranking of the experiments, $V_n(O, M_O)$ and $V_n(P, M_P)$.*

To understand this statement, recall that with rational agents, Blackwell's result B1 says that if O is more informative than P, then the agent is better off from O. For naive, the fundamental underlying experiment O can be more informative than P, yet may yield lower utility as he ultimately faces a different experiment due to his perception bias. However, even if we take into account what the naive's effective experiments are, OM_O and PM_P, it is still possible that the more informative experiment OM_O can have lower utility. This is due to the fact that the agent does not recognize the true information value of the experiments. It follows that Blackwell's main result on the comparison of experiments does not holds for biased agents.

Given the fact that there is no first period action in Blackwell, this result highlights when the worst-case dominance definition is trivially satisfied, and how first period actions' instrumental- and direct utility value can be separated.
First, suppose that the information experiment from x_1 is more informative than from y_1, $P_{x_1} \supset P_{y_1}$. If the direct utility from x_1 is lower than from y_1 at p, then worst-case dominance will always be satisfied by Result B3. However, if the utility from x_1 is higher at p, then we need to check whether worst-case-dominance is satisfied.
Similarly, when x_1 is less informative, $P_{x_1} \subset P_{y_1}$, but it is precisely chosen for it's direct utility, we again need to check worst-case dominance.

To understand the role of naivety, we present the Blackwell result for sophisticated agents:

Result B4: *: Take two experiments O and P, with perception biases M_O, M_P. $OM_O \supset PM_P$ if and only if $V_s(O, M_O) \geq V_s(P, M_P)$.*

The result follows from the observation that the sophisticated but biased agent's utility of O is equal to the rational agent's utility, who faces OM_O: $V_s(O, M_O) = V_s(OM_O, I) = V(OM_O)$. Clearly, this statement has no implication for the ordering of O and P itself. Observing an agent who takes an 'objectively' less informative experiments does therefore not imply that he is making an incorrect choice.[8] But this also implies that any sophisticated agent cannot be better off from an incorrect prior. The incorrect prior may cause him to choose an alternative action y_1 with a different underlying experiment. Since he correctly accounted for the lower-information value arising from his perception bias for any information experiment, this simply makes him choose worse first period actions.

Adding a Biased Perception to a Biased Priors

We now turn the previous analysis around. Taking the agent's biased prior as given, we show under which conditions adding or increasing a perception bias can make him better off. This result turns out to be more clear-cut. Adding any type of perception bias makes an agent worse-off if his chosen action profile is not 'too suboptimal' given the true p. Otherwise, there exists some type of perception bias that can improve the agent's welfare.

Proposition 3: *Let $\boldsymbol{x} = (x_1, \{x_a, x_b\})$ be a signal-sensitive action profile chosen at $\mu \neq p$.*
(a) If $E[u(\boldsymbol{x})|p] > \max\{E[u(x_1, x_a)|p], E[u(x_1, x_b)|p]\}$ then increasing the degree of perception bias strictly decreases expected utility for any type of perception bias.
(b) If $E[u(\boldsymbol{x})|p] < \max\{E[u(x_1, x_a)|p], E[u(x_1, x_b)|p]\}$, then increasing the degree of perception bias strictly increases expected utility for some type of bias.

Statement (a) says that the signal-sensitive **x** is preferred by an agent with the correct prior p over choosing the simple profile where the agent always chooses x_a or x_b in the second period. Adding noise to the signals is equivalent to putting more weight on a simple profile and strictly lowers μ's payoffs. This is exactly what any distortion function does.

Statement (b) states that one of the simple action profiles is preferred to the contingent one by the agent with belief p. In that case, the agent is strictly better off the more often he receives the respective signal realization that makes him choose such a profile.

We can also interpret Proposition 3 in terms of how severe the bias in prior is. To see this, recall that the action profile is signal sensitive at μ for exactly the reason that the posteriors following the signal fall into different parts of the partition that prescribes optimal second-period behavior. If we push up the prior, then after any signal action x_a is relatively more appealing compared to x_b than before. Pushed far enough the agent only wants to take x_a regardless of the signal realization; x_a is, after all, strictly

[8] To see this take a fully informative $n \times n$ vector P and some informative O of equal dimension. Let $M_P = [1/n]_{n\times n}$ i.e. the matrix filled only with $1/n$ which will turn P completely uninformative and keep the information value of O by setting $M_O = I$, so that $O \subset P$ but $OM_O \supset PM_P$.

better in state A. At this point, a perception bias can help the agent (b). If the prior is too close to p he is always worse off due to losing information.
Finally, it should also be clear form this discussion that for every signal sensitive profile, there always exists a region for either case.

Martingale bias

In the previous sections, the bias could generate a drift in beliefs whenever the perception bias was unbalanced between signals. A rational observer aware of the bias would form an ex-ante expectation over the posterior different from the prior. We now restrict the perception bias to what we call a 'martingale' bias: a bias such that the agent as well as the observer form the same expectation over posteriors. The martingale bias represents a simple random error in perception, that occurs indiscriminately of the signal type.[9] This section will help to qualify our previous Propositions. In contrast to Proposition 3, we show that under such a mistake, the agent with a biased prior can never be better off when a martingale perception bias is added. Changing the prior for an agent with such a perception bias can still be beneficial, however.

Consider a naive agent that thinks the signal strength is the true undistorted $\pi(x_1,\omega)$. Let $k_a(x_1,\mu) = k_b(x_1,\mu) \equiv \kappa \in [\frac{1}{2\pi(x_1,A)}, 1]$. We can write the agent's expected utility as

$$\begin{aligned} \mu \quad & [\kappa\pi(x_1,A)u(x_1,x_a|A) + (1-\kappa\pi(x_1,A))u(x_1,x_b|A)] \\ +(1-\mu) \quad & [(1-\kappa(1-\pi(x_1,B)))u(x_1,x_a|B) + \kappa(1-\pi(x_1,B))u(x_1,x_b|B)] \end{aligned} \tag{3.6}$$

From state independence, we know that $\kappa \leq 1$. Since the martingale bias represents a pure random error, in the extreme the signal should be pure noise and thus $\kappa \geq \frac{1}{2\pi(x_1,A)}$. As can be seen from Equation (3.6), if the true probability of receiving an a signal in state A is $\kappa\pi(x_1,A)$, but the agent believes this to be $\pi(x_1,A)$, there is too much weight on $u(x_1,x_a|A)$ and $u(x_1,x_b|B)$. The agent expects to achieve the high outcomes too often which overstates his perceived expected utility of the action profile. It follows that as the perception bias becomes worse, i.e. κ decreases, the agent is worse off. Moreover, this is true for any true p and any prior belief μ, so that there is no way that adding a martingale perception bias to a bias in prior is welfare improving:

Corollary 9: *For any $\kappa \in [\frac{1}{2\pi(x_1,A)}, 1]$ and any signal sensitive action profile $\{x_1,\{x_a,x_b\}\}$ chosen at μ, a decrease in κ decreases the expected utility for any μ and $p \in (0,1)$.*

[9] i.e. the garbling matrix is symmetric $M = \begin{bmatrix} m & 1-m \\ 1-m & m \end{bmatrix}$

Hence the restriction to a martingale bias yields a stronger result than Proposition 3. The martingale bias adds noise and thus reduces the signal strength equally for all signal types. Independent of the true state, this leads to strictly worse period 2 choices.

However, it is still possible that under such random error perception bias, the agent can benefit from an incorrect prior:

Corollary 10: *For any signal sensitive action profile* $\mathbf{x}$ *that is chosen at* μ*, and any action profile* $\mathbf{y} \neq \mathbf{x}$ *that is chosen at* $\mu' \neq \mu$*, there exists a* κ *such that* $E[u(\mathbf{y})] > E_d[u(\mathbf{x})]$ *when* $p = \mu \neq \mu'$ *if and only if*

$$\frac{\mu}{2}\left[u(x_1,x_a)|A) + u(x_1,x_b)|A)\right] + \frac{1-\mu}{2}\left[u(x_1,x_a)|B) + u(x_1,x_b)|B)\right] < E[u(\mathbf{y})|p] \tag{3.7}$$

Like Proposition 1 this assumes that the bias applies only to the chosen profile $\mathbf{x}$. We can also apply the bias symmetrically as long as we respect the condition that we can at most generate pure noise (rather than achieve a negative correlation). This means the experiment with the less precise signals determines the limit up to which we can garble symmetrically. We then get the following result:

Corollary 11: *For any signal sensitive action profile* $\mathbf{x}$ *that is chosen at* μ*, and any signal-sensitive action profile* $\mathbf{y} \neq \mathbf{x}$ *that is chosen at* $\mu' \neq \mu$ *and for* $\kappa_{min} \equiv \max\{\frac{1}{2\pi(x_1,A)}, \frac{1}{2\pi(y_1,A)}\}$*, there exists a* κ_{max} *with* $1 > \kappa_{max} > \kappa_{min}$ *such that* $E_d[u(\mathbf{y})] > E_d[u(\mathbf{x})]$ *when* $p = \mu \neq \mu'$ *for all* $\kappa \in [\kappa_{min}, \kappa_{max})$ *if and only if*

$$E[u(x)|p,\kappa_{min}] < E[u(y)|p,\kappa_{min}] \tag{3.8}$$

In other words, if we want to check whether an agent with prior μ' could be better off than one at the true p, we have to compare $\mathbf{y}$ to $\mathbf{x}$ at the extreme point where one of the two information experiments yields just noise. If at this point $\mathbf{y}$ is indeed preferred, such a reversal in the ordering of $\mathbf{x}$ and $\mathbf{y}$ is possible. But of course, this could already happen at a less extreme distortion.

In this section, and indeed in all previous ones, the agent is overconfident in the signal as he fails to discount the perception bias. The martingale property makes this quite similiar to a related major updating mistake found in the literature: in inference problems agents tend to sometimes over- and sometimes under-infer from a given sample of signals. For instance in the classical paper by Griffin and Tversky (1992) agents over-infer from small sample sizes but under-infer from large ones.[10] Indeed one can find a corresponding martingale bias setup for someone who over-infers: The two agents perceive the signal to be of the same strength and will receive the same signal with the same likelihood. However, the key difference between the two mistakes is that over-/ under-inference takes the true signal process as given and only changes the perception of signal strengths, whereas the distorted transmissions fixes the perception of signal quality but alters the actual probabilities with which signals are perceived. By changing the perceived signal strength, anything can happen in terms of welfare consequences. This lack of structure leads us to relegate the discussion to a footnote.[11]

3.6.3 Sophisticated vs Naive Agents

When the prior is correct, we have seen that a sophisticated agent is always weakly better off than a naive agent. We finish this section with a short example that highlights that this no longer needs to be true when μ differs from p. A sophisticated agent is aware of the reduced information value of each experiment and thus may want to take a different choice than a naive agent. If the true ordering of experiments under μ is different than under p, such behavioral adjustments can be suboptimal.

Example 4: Example 1 serves again as the basis. Suppose that the agent suffers from a martingale bias and tends to misperceive both a and b signals with probability $1-\kappa = 0.7$. Suppose further that $\mu = 0.57$.

Compared to the naive agent, the sophisticated is aware of the lower information value

[10] For a current, extensive summary of the experimental evidence on this topic, see Appendix B of Benjamin, Rabin and Raymond (2013).

[11] With over- or under-reaction to information, the agent can almost trivially achieve a higher utility regardless of which bias is adjusted:
First, suppose the sets of actions are $X_1 = \{x_1\}$ and $X_2 = \{x_a, x_b\}$. First, take the wrong prior as given: (1) Let the signal sensitive action profile $\mathbf{x}$ be optimal at p whereas the agent with prior $\mu \neq p$ prefers the simple profile (x_1, x_a). Then the agent is better off by exaggerating the signal quality so that his second period profile becomes signal sensitive. (2) Similarly, suppose the simple action profile (x_1, x_a) is optimal at p. Then an agent with $\mu \neq p$, who prefers the signal sensitive profile, can be better off by inferring too little if he prefers (x_1, x_a) over (x_1, x_b) in the absence of any signal.
Next take information processing bias as given: (3) Suppose that due to over-inference the agent prefers the conditional profile when it is optimal to choose (x_1, x_a). Changing the correct prior to some $\mu > p$ where the agent prefers (x_1, x_a) improves his utility.
The only non-trivial case (4) is under-inference and changing the correct prior. Suppose $\mathbf{x}$ is optimal at p but due to under-inference the agent with the correct prior prefers $\mathbf{z}$. Let $\mathbf{y}$ be the agent's preferred choice at some $\mu \neq p$. If the actual signal strength of $S(y_1)$ is sufficiently larger than of $S(z_1)$, then shifting the prior from p to μ helps the agent. To see why, notice due to under-inference his response to the signal is attenuated. However, the likelihood of receiving the correct signal under y_1 is higher (than he thinks) which will 'too often' select the 'correct' action from his conditional profile.

of x^I and opts for the simple profile (x^B, x^B). If $\mu = p$ he clearly does better than the naive who still chooses the information experiment in the first period. The naive agent, however, can do better than the sophisticated when the actual p is lower, take for example $p = 0.4$. His naivety causes him to pick the information experiment and subsequently to take x^A whenever he receives an a signal. For a low enough p this is always better than the sophisticated's choice of (x^B, x^B).

3.7 Example: Investment with Information Acquisition

We finish this paper with a simple investment problem that captures the fundamentals of the entrepreneur example outlined in the introduction.

An agent can invest his income w in either a risk-free or a risky asset, both having a unit price. There are two states of the world $\Omega = \{A, B\}$. The risk-free asset's return in any state of the world is normalised to zero. The risky asset yields a positive return of r_h in state A, but a negative return of r_l in state B. The agent has a prior $\mu_0 = \Pr(\omega = A)$ and is risk averse with $U(c) = \ln(c)$.

Before investing, the agent has the opportunity to purchase a signal $s \in \{a, b\}$ at a cost $c > 0$ that fully reveals ω. He then decides on the fraction of income he invests in the risky asset, denoted by x. Finally, let the agent's information perception bias be of the following structure: Whenever he receives a b signal - a signal that tells him not to invest - he might misinterpret it as an a signal with a probability δ. On the contrary, a signals are always perceived correctly.

We now show how holding an incorrect prior can help mitigate the bias in information perception through avoiding information.

First, we solve for x as a function of the posterior μ:

$$x(\mu, r_h, r_l) = \begin{cases} 0 & \text{if} \quad \mu r_h + (1-\mu) r_l \leq 0 \\ 1 & \text{if} \quad \mu \geq -r_l(r_h+1)/(r_h - r_l) \\ \frac{\mu r_h + (1-\mu) r_l}{-r_l r_h} & \text{otherwise} \end{cases}$$

Let $\Pr(\omega = A) = p = 0.55$. Let the risky asset's return be $r_h = 60\%$ in state A and $r_l = -95\%$ in state B. The agent's wealth is $w = 1.2$. The cost of acquiring information is $c = 0.25$. Furthermore, the agent's information bias is $\delta = 0.4$.

Notice that the expected return of the asset is $0.55 \cdot 0.6 - 0.45 \cdot 0.95 = -0.0975 < 0$. When $\mu = p$, an agent who does not acquire information would never invest in the risky asset.

If the agent acquires the fully informative signal, he will invest everything into the risky asset when he observes an a signal, and nothing at all following a b signal.
An unbiased agent (with the true prior) will acquire information since the information is cheap enough, $EU(\text{acquire info}) = p\ln((1+r_h)(w-c)) + (1-p)\ln(w-c) = 0.21 > \ln(1.2) = EU(\text{not acquire})$.
A naive agent with misperception bias makes the same choices as an unbiased one. When his prior is correct, his true utility from acquiring information is only

$$EU_{true}(\text{acquire info}|\mu = 0.55) = -0.33$$

since he sometimes invests his whole income when the state is B.[12]
If he instead holds a rather incorrect prior, i.e. $\mu = 0.9$, the information becomes of little value.[13] He consequently chooses not to acquire it and actually does better than if he had the correct prior.[14] While holding an initially more extreme belief, his posterior belief is lower than the naive agent who observed a (possibly wrong) a signal. This makes him invest only 78% of his total income, which is useful when the negative returns of an asset are large. On top of that, he benefits from saving on the information cost.

If we interpret the cost of acquiring information either as the direct cost of going to university or the opportunity cost of not working, this investment problem become analogous to the entrepreneur example that was outlined at the beginning. If the would-be entrepreneur holds a very inflated view about his abilities, he starts his business immediately (invests without information). If not, he goes to university. Suffering from a perception bias, he might then become even more convinced about himself, leading him to gamble even more on his endeavour. In expectation, the possibly useful information hurts him if the probability of misinterpreting it is too high.

[12] This is true for any $\mu \in [0.47, 0.89]$
[13] The results hold for any extreme prior in the interval $[0.89, 0.925]$. An even higher prior causes too much investment, lowering utility.
[14] His perceived expected utilities are: $EU_{perceived}(\text{acquire info}|\mu = 0.9) = -.39 < -0.37 = EU(\text{not acquire}|\mu = 0.9)$ while is actual utility (evaluated at the true prior) is $EU_{true}(\text{not acquire}|\mu = 0.9) = -0.22$.

3.8 Discussion

We have analyzed the interaction of two fundamental mistakes in information processing. We have seen how and when one bias can alleviate the problems caused by the other. The recent worrying rise of the anti-vaccination movement, with their tendency to disregard surmounting scientific evidence in favour of circumstantial observations that autism happens after (one of many) childhood infections, is just one example for the potential importance of these biases.

The idea that a mistake on one dimension may counteract some failure on another is clearly not new (Compte and Postlewaite (2004), De La Rosa (2011), etc.), yet to our knowledge it has never been explored in this basic updating setting. In contrast, the recent literature on confirmatory bias (Clements (2013), Roland G. Fryer, Harms and Jackson (2016)) as well as biases in information processing (Andreoni and Mylovanov (2012), Kartik, Lee and Suen (2016)) focuses solely on beliefs, how beliefs may converge or diverge in group, or whether information can be aggregated in social settings. It does not fully capture the dynamic aspect of payoff-relevant decisions as information experiments and their consequences on future payoffs. In that sense, we think of this paper as an extension of Blackwell (1951) to behavioral mistakes in information processing, for which we provide a full characterisation.

3.9 Appendix

3.9.1 A: Omitted Proofs

Proof of Result 2. Convexity follows directly from the fact that maximum expected utility is convex in μ. This result extends to two period model with intermediate signals.

Let $\vec{\mu}$ be row vector of belief about the different states Ω, $\vec{\mu} = (\Pr(\omega = A), \Pr(\omega = B))$. Define $\vec{u}(x)$ as the state contingent vector adjusted by signal probability and signal response:[15]

$$\vec{u}(x) = \begin{pmatrix} \pi(x_1, A)u(x_1, x_a|A) + (1-\pi(x_1, A))u(x_1, x_b|A) \\ \pi(x_1, B)u(x_1, x_a|B) + (1-\pi(x_1, B))u(x_1, x_b|B) \end{pmatrix}$$

Define maximum expected utility as a function of $\vec{\mu}$ as $g(\vec{\mu}) = \max_{x \in X} \vec{\mu} \cdot \vec{u}(x)$ and notice how this expression coincides with the rewritten expected utility formulation of equation 3.1.

To show $g(\mu)$ is convex, let x' (x'') optimal choice given the belief $\vec{\mu}'$ ($\vec{\mu}'$), and $\hat{x}$ the optimal choice for the convex combination $\vec{\hat{\mu}} = \lambda\vec{\mu}' + (1-\lambda)\vec{\mu}''$ with $\lambda \in (0,1)$. Then

$$\begin{aligned} g(\lambda\vec{\mu}' + (1-\lambda)\vec{\mu}'') &= (\lambda\vec{\mu}' + (1-\lambda)\vec{\mu}'') \cdot u(\hat{x}) \\ &\leq \lambda\vec{\mu}' \cdot u(u') + (1-\lambda)\vec{\mu}'' \cdot u(x'') \\ &= \lambda g(\vec{\mu}') + (1-\lambda) g(\vec{\mu}'') \end{aligned}$$

proving convexity. □

Proof of Result B2. By Blackwell's Result B1 it must be that $V(O) \geq V(OM)$. A rational (non-biased) agent would never choose the distorted over the undistorted experiment. A sophisticated agent is not able to choose the undistorted experiment, however. All he can do is to adjust his second period action profile anticipating the lower information value. His utility from O given M is equivalent to a rational agent facing the experiment OM. As a result the following inequalities hold: $V_s(O, M) = V_s(OM, I) = V(OM)$. This establishes the first inequality.

The naive agent acts as if he was fully rational, i.e. he applies the rationally optimal decision rule for P to OM. Consequently, he holds an incorrect posterior after observing the signal - which is equivalent to having wrong beliefs about the likelihood with

[15] compare this to the simple 1 period convexity proof, where $\vec{u}(x)$ is defined as the state contingent utility vector

which he receives a and b signals. He makes (weakly) worse decisions than a sophisticated agent; if not, the sophisticated agent could always imitate him. □

Proof of Result 4. Let the (unbiased) information experiment following action x_1 be of the form

$$P = \begin{bmatrix} \pi & 1-\pi \\ 1-\pi & \pi \end{bmatrix}$$

with $\pi = \pi(x_1, A)$. Let experiment P' be the experiment arising from the distortion function with the same action x_1, namely

$$P' = \begin{bmatrix} d_a(\pi, x_1, \mu) & d_b(1-\pi, x_1, \mu) \\ d_a(1-\pi, x_1, \mu) & d_b(\pi, x_1, \mu) \end{bmatrix}$$

If follows that $P \supset P^{'}$, as we can find an M such that $PM = P^{'}$. The corresponding M is:

$$M = \begin{bmatrix} k_a(x_1, \mu) & 1-k_a(x_1, \mu) \\ 1-k_b(x_1, \mu) & k_b(x_1, \mu) \end{bmatrix}$$

where $d_a(\pi, x_1, \mu) = k_a(x_1, \mu)\pi + (1 - k_b(x_1, \mu))(1-\pi)$ etc.

According to Result B1, a rational (non-biased) agent would never choose the distorted over the undistorted experiment. The utility arising from the distorted experiment is (weakly) lower.

To show the effect of an increase in the magnitude of the bias, we write the expected utility from x_1 as

$$p\left[\left(k_a(x_1,\mu)\pi + (1-k_b(x_1,\mu))(1-\pi)\right)u(x_1,x_a|A) + \left(k_b(x_1,\mu)(1-\pi) + (1-k_a(x_1,\mu))\pi\right)u(x_1,x_b|A)\right]$$
$$+(1-p)\left[\left(k_a(x_1,\mu)(1-\pi) + (1-k_b(x_1,\mu))\pi\right)u(x_1,x_a|B) + \left(k_b(x_1,\mu)\pi + (1-k_a(x_1,\mu))(1-\pi)\right)u(x_1,x_b|B)\right]$$

Suppose that initially k_b decreased. If we decrease it further to $k_b' < k_b$, the linearity property of the expected utility in probabilities implies that the expected utility must also decrease.

The same argument holds for any further decrease in the alternative parameter k_a.

Since there are no cross-interactions between changes of k_a and k_b, the same holds true for any joint change. □

Proof of Result 5. Note that we can write the matrix of the signal probabilities as a garbling of the matrix containing the true signal probabilities:

$$\begin{bmatrix} d_a(\pi, x_1, \mu) & d_b(1-\pi, x_1, \mu) \\ d_a(1-\pi, x_1, \mu) & d_b(\pi, x_1, \mu) \end{bmatrix} = \begin{bmatrix} \pi & 1-\pi \\ 1-\pi & \pi \end{bmatrix} \begin{bmatrix} k_a & 1-k_a \\ 1-k_b & k_b \end{bmatrix} \tag{3.9}$$

where the matrix on the right is the garbling matrix. Since the garbling matrix needs to be a Markov matrix, we know that $k_a, k_b \in [0,1]$.

Now suppose we increase the magnitude of the bias resulting in a distortion function d' and, contrary to the statement in the result, we have $d'_a(\pi, x_1, \mu) - d'_a(1-\pi, x_1, \mu) \geq d_a(\pi, x_1, \mu) - d_a(1-\pi, x_1, \mu)$. This implies that:

$$\pi k'_a + (1-\pi)(1-k'_b) - (1-\pi)k'_a - \pi(1-k'_b) \geq \pi k_a + (1-\pi)(1-k_b) - (1-\pi)k_a - \pi(1-k_b)$$

where k'_a, k'_b are the elements in the modified garbling matrix reflecting the increase in the magnitude of the bias. This can be rearranged to:

$$(2\pi - 1)(k'_a - (1-k'_b)) \geq (2\pi - 1)(k_a - (1-k_b))$$

But this requires that since $\pi > 1/2$, either $k'_a > k_a$, or $k'_b > k_b$, or both which contradicts the increase in the magnitude of the bias. □

Proof of Result 6. Take the action profile $(x_1, \{x_a, x_b\})$. It's actual expected utility is

$$\begin{aligned} p \quad & [\pi(x_1, A)u(x_1, x_a|A) + (1-\pi(x_1, A))u(x_1, x_b|A)] \\ +(1-p) \quad & [\pi(x_1, B)u(x_1, x_a|B) + (1-\pi(x_1, B))u(x_1, x_b|B)] \end{aligned}$$

If we evaluate this with some $\mu > p$, then the weight on the outcome in state A increases. This has two potential effects. Firstly, according to Result 12, a different action might be chosen in period 2 as the potential posteriors are now different. Furthermore, as follows from Result 1, a different x_1 might become optimal. Denote the potentially altered choice by $(y_1, \{y_a, y_b\})$. Any such alternative choice has the property that:

$$\begin{aligned} & [\pi(y_1, A)u(y_1, y_a|A) + (1-\pi(y_1, A))u(y_1, y_b|A)] \\ > & [\pi(x_1, A)u(x_1, x_a|A) + (1-\pi(x_1, A))u(x_1, x_b|A)] \end{aligned}$$

and

$$[\pi(y_1,B)u(y_1,y_a|B)+(1-\pi(y_1,B))u(y_1,y_b|B)]$$
$$<[\pi(x_1,B)u(x_1,x_a|B)+(1-\pi(x_1,B))u(x_1,x_b|B)]$$

As otherwise either the choice $(y_1,\{y_a,y_b\})$ would not be optimal as it yields lower expected utility when evaluated by μ or the choice under p would not be optimal since it results in lower utility in both states. It then immediately follows that if any choice of action is different when comparing optimal choices under μ and p, such a μ leads to a utility loss because of suboptimal choices. Furthermore, fixing some $\mu > p$ and iterating this argument for $\mu' > \mu$ yields the result. The equivalent argument applies to $\mu < p$. □

Proof of Result 7. If the agent prefers $\mathbf{y}$ over $\mathbf{x}$, i.e. $E[u(\mathbf{x})] < E[u(\mathbf{y})]$, the statement is trivially true. If it does not hold, we show that there is such a d that makes

$$E_d[u(\mathbf{x})|\mu] < E_d[u(\mathbf{y})|\mu]$$

According to $D1$, we have that $\mu u(x_1,x_s|A)+(1-\mu)u(x_1,x_s|B) < E[u(\mathbf{y})|\mu]$ for some $s \in \{a(x_1),b(x_1)\}$.

Suppose wlog that $s = a(x_1)$. This implies that small enough $\epsilon > 0$, we have

$$\begin{aligned}&(1-\epsilon)\left[\mu u(x_1,x_a|A)+(1-\mu)u(x_1,x_a|B)\right]+\epsilon\left[\mu u(x_1,x_b|A)+(1-\mu)u(x_1,x_b|B)\right]\\&=\mu\left[(1-\epsilon)u(x_1,x_a|A)+\epsilon u(x_1,x_b|A)\right]+(1-\mu)\left[(1-\epsilon)u(x_1,x_a|B)+\epsilon u(x_1,x_b|B)\right]\\&<E[u(\mathbf{y})|\mu]\end{aligned} \tag{3.10}$$

which quite resembles the maximum expected utility representation (3.1).

Now set $k_a(x_1,\mu)=1$ and $k_b(x_1,\mu)=\frac{\epsilon}{\pi(x_1,A)}$. This causes the agent to receive the worst-case signal a in state B with probability $d_a(\pi(x_1|B),S(x_1),\mu)=1-\epsilon$. Clearly, this also makes him receive the (correct) worst-case signal a in state A with probability $d_a(\pi(x_1|A),S(x_1),\mu)>1-\epsilon$. Consequently, the actual expected utility following this distortion is slightly larger than the left-hand side of Equation 3.10 above, but as ϵ becomes sufficiently small, $E_d[u(\mathbf{x})|\mu] < E[u(\mathbf{y})|\mu]$. Finally let there be no distortion for the signals $S(y_1)$, i.e. $d_a(\pi,S(y_1),\mu)=\pi$ for all $\pi \in (0,1)$. The result follows. □

Proof of Proposition 2. Sufficiency is an almost immediate consequence of Result 7. We know there exists d, such that $E_d[u(\mathbf{y})|\mu] > E_d[u(\mathbf{x})|\mu]$ by worst-case dominance at μ and that an agent with belief μ' chooses $\mathbf{y}$ compared to $\mathbf{x}$. Finally simply set $\mu = p$.

In the other direction, suppose such an improvement is possible. Since any distortion garbles the information, $E[u(\mathbf{y})|p] \geq E_d[u(\mathbf{y})|p]$. But if worst-case dominance does not hold there is no marginal distribution over $S(x_1)$ such that $E_d[u(\mathbf{x})|p] < E[u(\mathbf{y})|p]$ since $E_d[u(\mathbf{x})|p] \geq \min_{s \in S(x_1)} E[u(x_1, x_s)] > E[u(\mathbf{y})|p]$. A contradiction. □

Proof of Corollary 8. Sufficiency: Similar to Proposition 2, consider the extreme distortion d that only sends the signal

$$\underline{s}_x = \underset{s}{\operatorname{argmin}}\, \mu u(x_1, x_s)|A) + (1-\mu)u(x_1, x_s|B)$$

$$\overline{s}_x = \underset{s}{\operatorname{argmax}}\, \mu u(x_1, x_s)|A) + (1-\mu)u(x_1, x_s|B)$$

If simple dominance holds, at least one of those signals must yield lower expected utility under x than under y.

Necessity: Suppose such a distortion exists. Then

$$E[u(y)|p] < E[u(x)|p]$$

while

$$E_d[u(y)|p] > E_d[u(x)|p]$$

We can write the actual expected utility explicitly as a linear combination between the a and the b signal.

$$\begin{aligned} p \quad & \pi(x_1, A)u(x_1, x_a|A) \quad + \quad (1-p)\pi(x_1, B)u(x_1, x_a|B) \\ +p \quad & (1-\pi(x_1, A))u(x_1, x_b|A) \quad + \quad (1-p)(1-\pi(x_1, B))u(x_1, x_b|B) \end{aligned}$$

and similarly for y. Any distortion can either lower terms or shift weight from one to the other. If the distortion is such that weight is added in the sense that both $\pi(x_1, A)$ and $\pi(x_1, B)$ are increased or decreased then if $E_d[u(y)|p] > E_d[u(x)|p]$ it must be that either

$$pu(x_1, x_a|A) + (1-p)u(x_1, x_a|B) < pu(y_1, y_a|A) + (1-p)u(y_1, y_a|B) \tag{3.11}$$

or

$$pu(x_1, x_b|A) + (1-p)u(x_1, x_b|B) < pu(y_1, y_b|A) + (1-p)u(y_1, y_b|B) \tag{3.12}$$

or both which guarantees simple. If the distortion lowers both terms i.e. by garbling both type of signals equally, then $\pi(x_1, B)$ and $(1-\pi(x_1, A))$ increase. If this increase leads to a reversal in the expected utility ranking of **x** and **y** in the sense that $E_d[u(\mathbf{y})|p] > E_d[u(\mathbf{x})|p]$, then because of linearity this must also be true for the most extreme case with $d_a(\pi(x_1, A)) = d_b(\pi(x_1, B)) = \frac{1}{2}$. But then

$$\frac{p}{2}u(x_1, x_a|A) + \frac{(1-p)}{2}u(x_1, x_a|B) < \frac{p}{2}u(y_1, y_a|A) + \frac{(1-p)}{2}u(y_1, y_a|B) \tag{3.13}$$

or

$$\frac{p}{2}u(x_1, x_b|A) + \frac{(1-p)}{2}u(x_1, x_b|B) < \frac{p}{2}u(y_1, y_b|A) + \frac{(1-p)}{2}u(y_1, y_b|B) \tag{3.14}$$

or both which means at least one of them is also true for $d_a(\pi(x_1, A)) = d_a(\pi(x_1, B)) = 1$ or $d_b(1-\pi(x_1, B)) = d_b(1-\pi(x_1, A)) = 1$ which again guarantees simple dominance. □

Proof of Result B3. To see take two separate experiments O and P, with $O \supset P$. First, let M_O be such that $OM_O = P$ and let P be undistorted, i.e. $M_P = I$. It follows that the naive agent is better off from P than O as $V_n(P, I) = V(P) > V_n(O, M_1)$ (the inequality is strict if the agent is signal sensitive). But by continuity this also implies that there exist some M_O, M_P such that $OM_O \supset PM_P$ yet $V_n(P, M_P) > V_n(O, M_O)$. As a result, the naive's expected utility can neither be ordered by the perceived information experiments, nor the actual ones. In the other direction, we know that for $V_n(P, M_P) > V_n(O, M_O)$ it is possible that $OM_O \supset PM_P$, as well as $O \supset P$. □

Proof of Proposition 3. Statement (a) says that the signal-sensitve **x** is preferred by an agent with the correct prior p over choosing the simple profile where the agent chooses either always x_a or always x_b in the second period. Adding noise weakens the correlation with the state and the dominated outcomes occur more frequently.Hence the agent with prior μ can never be better off.

Statement (b) highlights that one of the simple action profiles is preferred to the contingent one by the agent with belief p. Assume wlog that the preferred action profile is (x_1, x_a). Any distortion function that increases the likelihood of an a signal will make the agent better off as he will take action x_a as a result, i.e. set $k_a(x_1, \mu) = 1$

and $k_b(x_1,\mu) \in (0,1)$. Clearly, the agent's expected utility is increasing as $k_b(x_1,\mu)$ decreases to 0. The extreme $k_b(x_1,\mu) = 0$ represents a perception bias where all b signals are interpreted as a signals and the agent always takes the (second) best simple action profile. □

Proof of Corollary 9. A decrease in κ leads to a less informative signal.

$$\begin{aligned} \mu \quad & [\kappa\pi(x_1,A)u(x_1,x_a|A) + (1-\kappa\pi(x_1,A))u(x_1,x_b|A)] \\ +(1-\mu) \quad & [(1-\kappa(1-\pi(x_1,B)))u(x_1,x_a|B) + \kappa(1-\pi(x_1,B))u(x_1,x_b|B)] \end{aligned}$$

Trivially, the expected utility at μ as written above decreases as there is more weight on $u(x_1,x_b|A)$ and $u(x_1,x_a|B)$. The agent is unaware of this and his choice in unaffected by κ. For any κ, he evaluates the profile as follows:

$$\begin{aligned} \mu \quad & [\pi(x_1,A)u(x_1,x_a|A) + (1-\pi(x_1,A))u(x_1,x_b|A)] \\ +(1-\mu) \quad & [(1-\pi(x_1,B))u(x_1,x_a|B) + (1-\pi(x_1,B))u(x_1,x_b|B)] \end{aligned}$$

The true p could be different from μ, so we just need to show that for any p, a decrease in κ has a negative welfare effect. The actual expected utility is:

$$\begin{aligned} p \quad & [\kappa\pi(x_1,A)u(x_1,x_a|A) + (1-\kappa\pi(x_1,A))u(x_1,x_b|A)] \\ +(1-p) \quad & [(1-\kappa(1-\pi(x_1,B)))u(x_1,x_a|B) + \kappa(1-\pi(x_1,B))u(x_1,x_b|B)] \end{aligned}$$

Simply notice since $u(x_1,x_b|A) < (x_1,x_a|A)$, as well as $u(x_1,x_a|B) < u(x_1,x_b|B)$, the random noise indeed lowers expected utility for any p. □

Proof of Corollary 10. If: Suppose the inequality in Equation 3.7 holds. Then a distortion of $\kappa = \frac{1}{2\pi(x,A)}$ turns the signals into complete noise. If this applies to the profile $\mathbf{x}$ only, $\mathbf{y}$ yields higher expected utility at p even though without any distortion $E[\mathbf{x}|p] > E[\mathbf{y}|p]$.

Only if: Suppose the property does not hold. The for any distortion $\kappa \in [\frac{1}{2\pi(x,A)},1]$, the expected utility from $\mathbf{x}$ at p exceeds the one from $\mathbf{y}$ as for any such κ,

$$E_d[\mathbf{x}|p] \geq \frac{\mu}{2}[u(x_1,x_a)|A) + u(x_1,x_b)|A)] + \frac{1-\mu}{2}[u(x_1,x_a)|B) + u(x_1,x_b)|B)]$$

□

Proof of Corollary 11. If: Follows directly as it is the case for $\kappa = \kappa_{min}$ and continuity guarantees that for small enough $\epsilon > 0$ such that it is also the case for $\kappa = \kappa_{min} + \epsilon$

Only If: Note that for $\mathbf{x}$ to be chosen at p, we need that $E[\mathbf{x}|p] > E[\mathbf{y}|p]$ as the agent is unaware of the perception bias. The effect of a decrease in κ is

$$\begin{aligned} -\frac{\partial E_d[u(\mathbf{x})|p]}{\partial \kappa} = \quad & -p\pi(x_1, A)[u(x_1, x_a|A) - u(x_1, x_b|A)] \\ & -(1-p)\pi(x_1, A)[u(x_1, x_b|B) - u(x_1, x_a|B)] \end{aligned} \tag{3.15}$$

This does not depend on κ itself. Therefore, if

$$\left[\frac{\partial E_d[u(\mathbf{x})|p]}{\partial \kappa}\right]_{\kappa=1} > \left[\frac{\partial E_d[u(\mathbf{y})|p]}{\partial \kappa}\right]_{\kappa=1}$$

then the inequality also holds for all allowable κ. Take the $\kappa_e \in (\kappa_{min}, \kappa_{max})$ such that $E_d[\mathbf{x}|p, \kappa_e] = E_d[\mathbf{y}|p, \kappa_e]$. If this does not exist, then $\mathbf{y}$ is never the best choice which yields the result. If it does exist, then it must mean that

$$\left[\frac{\partial E_d[u(\mathbf{x})|p]}{\partial \kappa}\right]_{\kappa=1} > \left[\frac{\partial E_d[u(\mathbf{y})|p]}{\partial \kappa}\right]_{\kappa=1}$$

But then for all $\kappa \in [\kappa_{min}, \kappa_e)$, the inequality also holds and therefore

$$E[u(x)|p, \kappa_{min}] < E[u(y)|p, \kappa_{min}]$$

as desired. □

3.9.2 B: Relationship To Blackwell

This section provides a general introduction to Blackwell, going beyond the 2-stage, 2-signal case, and providing clear definitions of value functions. In Blackwell's comparison of experiments, an agent compares two statistical information experiments, P and Q, which generate a signal about the states of the world $\Omega = \{\omega_1, ..., \omega_n\}$. As before, an information experiment is defined by a $n \times n_K$ Markov matrix K, where element K_{ij} refers to the probability of receiving a signal $s = j$ in state ω_i. We denote the set of possible signals from K by $S^K = \{1, 2, \ldots, N_K\}$.

The agent is endowed with utility function $u(x, \theta_i)$, defined over $A \times \Omega$. After receiving a signal s^K, he chooses an action x from his finite action set X. His prior is denoted by the vector $\mu_0 = \{\mu(1), ..., \mu(n)\}$, whereas his posterior after receiving signal s from experiment K is denoted by $\mu(s^K)$. The agent maximizes expected utility:

$$\max_{x \in X} \sum_i \mu(i|s^K)\, u(x|\omega_i)$$

Denote the solution to this problem by $x^*(\mu(s^K))$ and the corresponding value function $V(\mu(s^K))$.
The agent's probability of receiving an s^P signal is $\pi(s^P) = \sum_{i \in n} P_{is^P}\mu_0(i)$. and expected utility from observing experiment P is $V(P) = \sum_{s^P \in S^P} \pi(s^P) V(\mu(S^P))$.

Definition 12. *Let P and Q be two experiments. We say that experiment P is more informative than Q, denoted by $P \supset Q$, if there is a $n_P \times n_Q$ Markov matrix M with $PM = Q$.*

This definition highlights how the information experiment P is garbled into Q by the matrix M; as if noise was added to create a less informative experiment. From this statistical notion, Blackwell derives his famous result:

Result B5: *Let P and Q be two $n \times n_P$ and $n \times n_Q$ Markov matrices. $V(P) \geq V(Q)$ if and only if there is a $n_P \times n_Q$ Markov matrix M such that $PM = Q$.*

Simply put, a more informative experiment implies a higher utility. Surprisingly, the converse is also true: If one experiment yields more expected utility than another, it must be that it is more informative.
In Blackwell's setup the agent is rational. He fully understands the signal generating structure of the two experiments and correctly perceives each signal. The garbling matrix doesn't represent the agent's misperception but is a tool to order the experiments by informativeness. Moreover, the agent does not take any first period choice. All that occurs is a welfare comparisons of the two experiments. Needless to say, if the rational agent had to make a choice in Blackwell, he would prefer the more informative experiment.
In order to talk about a Blackwell setting when agents misperceive information, we need a new function representing the agents' expected utility from the experiments. Suppose the unbiased, true information experiment is P but the agent misperceives some signals according to matrix M. As a result, the agent faces an information experiment PM. We denote the expected utility of a naive agent unaware of his perception bias by $V_n(P, M)$, where the first argument always indicates the original signal generating process and the second argument the agent's misperception. This agent believes the information experiment is P and thus chooses his action as if he was rational facing P. If the agent is sophisticated, that is he is cognisant of his misperception M, we write his utility as: $V_s(P, M) = \sum_{s^P \in S^{PM}} \pi(s^{PM}) V(\mu(S^{PM}))$.

3.9.3 C: Decision Problem Appendix

Result 12: *For any action* $x_1 \in X_1$ *and any* $x_2 \in X_2$ *such that for some* $\mu(s) \in (0,1)$ *we have* $E[u(x_1, x_2)|\mu(s)] \geq E[u(x_1, y)|\mu(s)]$ $\forall y \in X_2$, *there exists an interval* $\mathcal{I} \subset [0,1]$ *with the property that* x_2 *is chosen if and only if* $\mu(s) \in \mathcal{I}$.

Proof of Result 12. If for some $(x_1, x_2) \in X$ we have

$$E[u(x_1, x_2)|\mu(s)] \geq E[u(x_1, y_2)|\mu(s)] \quad \forall y_2 \in X_2$$

then continuity in $\mu(s)$ implies that this is also true for at least some $[\mu(s), \overline{\mu})$ or $(\underline{\mu}, \mu(s)]$. If this interval is $[0,1]$, the proof is complete. If not, there is some μ_j for which $x_j \in X_2$ is the optimal choice. Then either $u(x_1, x_j|A) > u(x_1, x_2|A)$ or $u(x_1, x_j|B) > u(x_1, x_2|B)$. This shows us that for either $\mu > \mu_j$ or $\mu < \mu_j$, x_j achieves a higher expected payoff than x_2. This completes the proof. □

Bibliography

Alesina, Alberto, Rafael di Tella, and Robert MacCulloch. 2004. "Inequality and happiness: are Europeans and Americans different?" *Journal of Public Economics*, 88.

Andreoni, James, and Tymofiy Mylovanov. 2012. "Diverging Opinions." *American Economic Journal: Microeconomics*, 4(1): 209–232.

Ashraf, Nava, Oriana Bandiera, and Scott S. Lee. 2014. "Awards Unbundled: Evidence from a Natural Field Experiment." *Journal of Economic Behavior & Organization*, 100.

Aumann, Robert J. 1987. "Correlated equilibrium as an expression of Bayesian rationality." *Econometrica*, 1–18.

Aumann, Robert J., and Jacques H. Dreze. 1974. "Cooperative games with coalition structures." *International Journal of Game Theory*, 3(4): 217–237.

Barrow, Lisa. 2002. "School choice through relocation: evidence from the Washington, D.C. area." *Journal of Public Economics*, 86(2): 155 – 189.

Baumann, Leonie. 2015. "Time Allocation in Friendship Networks." University of Cambridge - Cambridge-INET Institute.

Bayer, Patrick, Fernando Ferreira, and Robert McMillan. 2007. "A Unified Framework for Measuring Preference for Schools and Neighborhoods." *Journal of Political Economy*, 115(4).

Becker, Gary S. 1973. "A Theory of Marriage: Part I." *Journal of Political Economy*, 81(24): 813 – 846.

Becker, Gary S., Kevin M. Murphy, and Iván Werning. 2005. "The Equilibrium Distribution of Income and the Market for Status." *Journal of Political Economy*, 113(2): 282–310.

Benabou, Roland, and Jean Tirole. 2002. "Self- Confidence and Personal Motivation." *Quarterly Journal of Economics*, 117(3): 817–915.

Benjamin, Daniel J., Matthew Rabin, and Collin Raymond. 2013. “A Model of Non-Belief in the Law of Large Numbers.” working paper.

Benoit, JeanPierre, and Don A. Moore. 2015. “Does the Better-Than-Average Effect Show That People Are Overconfident? Two Experiments.” *Journal of the European Economic Association*, 13(2): 293–329.

Benoit, JeanPierre, and Juan Dubra. 2011. “Apparent Overconfidence.” *Econometrica*, 79(5): 1591–1625.

Blackwell, David. 1951. “Comparison of experiments.” 93–102. University of California Press.

Blackwell, David, and M. A. Girshick. 1954. *Theory of Games and Statistical Decisions.* New York:John Wiley & Sons.

Bloch, Francis, and Bhaskar Dutta. 2009. “Communication networks with endogenous link strength.” *Games and Economic Behavior*, 66(1): 39 – 56.

Board, Simon. 2009. “Monopolistic group design with peer effects.” *Theoretical Economics*, 4(1): 89 – 125.

Bottan, Nicolas L., and Ricardo Perez-Truglia. 2017. “Choosing Your Pond: Location Choices and Relative Income.” National Bureau of Economic Research Working Paper 23615.

Bramoullé, Yann, and Rachel Kranton. 2007. “Public goods in networks.” *Journal of Economic Theory*, 135(1): 478 – 494.

Brito, Dagobert L, and William H Oakland. 1980. “On the monopolistic provision of excludable public goods.” *The American Economic Review*, 70(4): 691–704.

Brown, Gordon, Jonathan Gardner, Andrew Oswald, and Jing Qian. 2008. “Does wage rank affect employees’ well-being.” *Industrial Relations*, 47.

Bruner, Jerome S., and Mary C. Potter. 1964. “Interference in Visual Recognition.” *Science*, 144(3617): 424–425.

Brunnermeier, Markus K., and Jonathan A. Parker. 2005. “Optimal Expectations.” *The American Economic Review*, 95(4): 1092–1118.

Buchanan, James M. 1965. “An Economic Theory of Clubs.” *Economica*, 32(125): 1–14.

Burks, Stephen V., Jeffrey P. Carpenter, Lorenz Goette, and Aldo Rustichini. 2013. "Overconfidence and Social Signalling." *The Review of Economic Studies*, 80(3): 949–983.

Cabrales, Antonio, Antoni Calvó-Armengol, and Yves Zenou. 2011. "Social interactions and spillovers." *Games and Economic Behavior*, 72(2): 339 – 360.

Calabrese, Stephen, Dennis Epple, Thomas Romer, and Holger Sieg. 2006. "Local Public Good Provision: Voting, peer effects and mobility." *Journal of Public Economics*, 90.

Campbell, Karen E., Peter V. Marsden, and Jeanne S. Hurlbert. 1986. "Social resources and socioeconomic status." *Social Networks*, 8(1): 97 – 117.

Card, David, Alexandre Mas, Enrico Moretti, and Emmanuel Saez. 2012. "Inequality at Work: The Effect of Peer Salaries on Job Satisfaction." *American Economic Review*, 102.

Carrillo, Juan D., and Thomas Mariotti. 2000. "Strategic Ignorance as a Self-Disciplining Device." *The Review of Economic Studies*, 67(3): 529–544.

Charness, Gary, Aldo Rustichini, and Jeroen van de Ven. 2014. "Self-Confidence and Strategic Behavior." working paper.

Clements, Matthew T. 2013. "Is more information better? Social learning with confirmatory bias." *working paper*.

Compte, Oliver, and Andrew Postlewaite. 2004. "Confidence-Enhanced Performance." *The American Economic Review*, 94(5): 1536–1557.

Conley, John P., and Myrna H. Wooders. 2001. "Tiebout Economies with Differential Genetic Types and Endogenously Chosen Crowding Characteristics." *Journal of Economic Theory*, 98(2): 261–294.

Darley, John M., and Paget H. Gross. 1983. "A Hypothesis-Confirming Bias in Labeling Effects." *Journal of Personality and Social Psychology*, 44(1): 20–33.

De La Rosa, Leonidas Enrique. 2011. "Overconfidence and moral hazard." *Games and Economic Behavior*, 73: 429–451.

Demange, Gabrielle, and Dominique Henriet. 1991. "Sustainable oligopolies." *Journal of Economic Theory*, 54(2): 417 – 428.

Duesenberry, James. 1949. *Income, Saving and the Theory of Consumer Behaviour.* Harvard University Press.

Durieu, Jacques, Hans Haller, and Philippe Solal. 2011. "Nonspecific Networking." *Games*, 2(1).

Edlin, Aaron S, and Chris Shannon. 1998. "Strict single crossing and the strict Spence-Mirrlees condition: a comment on monotone comparative statics." *Econometrica*, 66(6): 1417–1425.

Epple, Dennis, and Richard E. Romano. 2011. "Chapter 20 - Peer Effects in Education: A Survey of the Theory and Evidence." In . Vol. 1 of *Handbook of Social Economics*, , ed. Jess Benhabib, Alberto Bisin and Matthew O. Jackson, 1053 – 1163.

Epple, Dennis, Thomas Romer, and Holger Sieg. 2001. "Interjurisdictional Sorting and Majority Rule: An Empirical Analysis." *Econometrica*, 69(6).

Fischoff, Baruch, Paul Slovic, and Sarah Lichtenstein. 1977. "Knowing with Certainty: The Appropriateness of Extreme Confidence." *Journal of Experimental Psychology: Human Perception and Performance*, 3(4): 552–564.

Frank, Robert H. 1985. *Choosing the Right Pond.* Oxford University Press.

Frank, Robert H. 2005. "Positional Externalities Cause Large and Preventable Welfare Losses." *The American Economic Review*, 95(2): 137–141.

Ghiglino, Christian, and Antonella Nocco. 2017. "When Veblen meets Krugman: social network and city dynamics." *Economic Theory*, 63(2): 431–470.

Ghiglino, Christian, and Sanjeev Goyal. 2010. "Keeping up with the neighbors: social interaction in a market economy." *Journal of the European Economic Association*, 8(1): 90–119.

Griffin, Dale, and Amos Tversky. 1992. "The Weighing of Evidence and the Determinants of Confidence." *Cognitive Psychology*, 24: 411–35.

Guthrie, Chris, Jeffrey Rachlinski, and Andrew Wistrich. 2001. "Inside the Judicial Mind: Heuristics and Biases." *Cornell Law Review*, 56: 777–830.

Heffetz, Ori, and Robert H. Frank. 2011. "Chapter 3 - Preferences for Status: Evidence and Economic Implications." In . Vol. 1 of *Handbook of Social Economics*, , ed. Jess Benhabib, Alberto Bisin and Matthew O. Jackson, 69 – 91.

He, Wei, Xiang Sun, and Yeneng Sun. 2017. "Modeling infinitely many agents." *Theoretical Economics*, 12.

Hirshleifer, Jack. 1971. "The Private and Social Value of Information and the Reward to Inventive Activity." *The American Economic Review*, 61(4): 561–574.

Hopkins, Ed, and Tatiana Kornienko. 2004. "Running to Keep in the Same Place: Consumer Choice as a Game of Status." *American Economic Review*, 94.

Jemmott, John B, and Elida Gonzalez. 1989. "Social status, the status distribution, and performance in small groups." *Journal of Applied Social Psychology*, 19(7): 584 – 598.

Kacperczyk, Aleksandra J. 2013. "Social Influence and Entrepreneurship: The Effect of University Peers on Entrepreneurial Entry." *Organization Science*, 24(3): 664–683.

Kartik, Navin, Frances Xu Lee, and Wing Suen. 2016. "A Theorem on Bayesian Updating and Applications to Signaling Games." *working paper*.

Kendall, Todd D. 2003. "Spillovers, Complementarities, and Sorting in Labor Markets with an Application to Professional Sports." *Southern Economic Journal*, 70(2).

Landier, Augustin, and David Thesmar. 2009. "Financial Contracting with Optimistic Entrepreneurs." *Review of Financial Studies*, 22(1): 177–150.

Levy, Gilat, and Ronny Razin. 2015. "Preferences over Equality in the Presence of Costly Income Sorting." *American Economic Journal: Microeconomics*, 7(2): 308 – 337.

Lichtenstein, Sarah, Baruch Fischoff, and Phillips Lawrence. 1982. "Calibration of Probabilities: The State of the Art to 1980." In *Judgement under uncertainty: Heuristics and biases.* , ed. Daniel Kahneman, Paul Slovic and Amos Tversky, 306–334. Cambridge:Cambridge University Press.

Ludwig, Sandra, Philipp C. Wichardt, and Hanke Wickhorst. 2011. "Overconfidence can improve an agent's relative and absolute performance in contests." *Economic Letters*, 110: 193–196.

Luttmer, Erzo. 2005. "Neighbors as Negatives: Relative Earnings and Well-Being." *Quarterly Journal of Economics*, 120.

Maccheroni, Fabio, Massimo Marinacci, and Aldo Rustichini. 2012. "Social Decision Theory: Choosing Within and Between Groups." *Review of Economic Studies*, 79(4).

Malmendier, Ulrike, and Geoffrey Tate. 2005. "CEO Overconfidence and Corporate Investment." *Journal of Finance*, 60(6): 2661–2700.

Marschak, Jacob, and Koichi Miyasawa. 1968. "Economic Comparability of Information Systems." *International Economic Review*, 9(2): 137–174.

Mas, Alexandre, and Enrico Moretti. 2009. "Peers at Work." *American Economic Review*, 99(1).

Moldovanu, Benny, Aner Sela, and Xianwen Shi. 2007. "Contests for status." *Journal of Political Economy*, 115(2): 338–363.

Mussa, Michael, and Sherwin Rosen. 1978. "Monopoly and product quality." *Journal of Economic Theory*, 18(2): 301 – 317.

Myerson, Roger B. 1982. "Optimal coordination mechanisms in generalized principal–agent problems." *Journal of mathematical economics*, 10(1): 67–81.

Pack, Howard, and Janet Rothenberg Pack. 1977. "Metropolitan Fragmentation and Suburban Homogeneity." *Urban Studies*.

Persky, Joseph. 1990. "Suburban Income Inequality." *Regional Science and Urban Economics*.

Rabin, Mathew, and Joel Schrag. 1999. "First Impressions Matter: A Model of Confirmatory Bias." *Quarterly Journal of Economics*, 114(2): 37–82.

Ray, Debraj, and Arthur Robson. 2012. "Status, Intertemporal Choice, And Risk-Taking." *Econometrica*, 80.

Rayo, Luis. 2013. "Monopolistic Signal Provision." *B.E. Journal of Theoretical Economics*, 13(1): 27 – 58.

Robson, Arthur. 1992. "Status, the Distribution of Wealth, Private and Social Attitudes to Risk." *Econometrica*, 60.

Roland G. Fryer, Jr., Philipp Harms, and Matthew O. Jackson. 2016. "Updating Beliefs when Evidence is Open to Interpretation: Implications for Bias and Polarization." *working paper*.

Scotchmer, Suzanne. 2005. "Consumption Externalities, Rental Markets and Purchase Clubs." *Economic Theory*, 25.

Sobel, Joel. 2002. "Can We Trust Social Capital?" *Journal of Economic Literature*, 40.

Stark, Oded, and J. Edward Taylor. 1991. "Migration incentives, migration types: the role of relative deprivation." *Economic Journal*, 101: 1163 – 1178.

Tiebout, Charles M. 1956. "A Pure Theory of Local Expenditures." *Journal of Political Economy*, 64(5): 416–424.

Veblen, Thorstein. 2007. *The theory of the leisure class. Oxford world's classics,* Oxford University Press.

Wallace, Amy. 2009. "How Panicked Parents Skipping Shots Endanger Us All." `https://www.wired.com/2009/10/ff_waronscience`, Accessed: 2017-10-20.

Weinstein, Neil D. 1980. "Unrealistic Optimism about Future Life Events." *Journal of Personality and Social Psychology*, 39(5): 806–820.

Weiss, Yoram, and Chaim Fershtman. 1998. "Social status and economic performance: A survey." *European Economic Review*, 42: 801 – 820.

Windsteiger, Lisa. 2017. "Essays on sorting and inequality." PhD diss. London School of Economics and Political Science.